

ABSTRACT

Epistemic Blame: Its Nature and its Norms

James C. Bryant, Ph.D.

Mentor: Trent G. Dougherty, Ph.D.

In this dissertation, I investigate our practices of blaming others for failing to believe as they ought to believe. I begin by articulating an account of blame in general, and extend that account to specifically epistemic blame. After considering the relationship between epistemic blame and moral blame, I argue that it is very difficult for us to know whether others are epistemically blameworthy. I conclude by arguing that we have good reasons to expect genuine epistemic blameworthiness to be quite rare, and that this fact justifies a charitable reluctance to blame others epistemically.

Epistemic Blame: Its Nature and its Norms

by

James C. Bryant, B.A., M.A.

A Dissertation

Approved by the Department of Philosophy

Michael D. Beaty, Ph.D., Chairperson

Submitted to the Graduate Faculty of
Baylor University in Partial Fulfillment of the
Requirements for the Degree
of
Doctor of Philosophy

Approved by the Dissertation Committee

Trent G. Dougherty, Ph.D., Chairperson

Michael D. Beaty, Ph.D.

Alexander R. Pruss, Ph.D.

Francis J. Beckwith, Ph.D.

Jonathan Tran, Ph.D.

Accepted by the Graduate School

August 2016

J. Larry Lyon, Ph.D., Dean

Copyright © 2016 by James C. Bryant

All rights reserved

TABLE OF CONTENTS

ACKNOWLEDGMENTS	v
PREFACE	vi
CHAPTER ONE: A General Account of Blame	1
Introduction	1
Competing Accounts of Blame	2
Arguments Against Cognitive Accounts	13
For The Account	25
CHAPTER TWO: What is Epistemic Blame?	28
Introduction	28
Norms and Goals	29
Epistemic Norms and the Epistemic Goal	31
The Epistemic Norms	35
The Other Norms	38
CHAPTER THREE: Epistemic and Moral Norms	50
Introduction	50
Clifford and Locke	51
Arguments Against Identifying Epistemic and Moral Norms	60
Why Think that Epistemic Blame is Moral Blame?	70
CHAPTER FOUR: Blaming Well is Hard	75
Introduction	75
The Evidentialist Norm	75
The Responsibilist Norm	88
CHAPTER FIVE: Epistemic Blame and Charity	91
Introduction	91
Why be Charitable?	92
Why not Withhold?	96
Principles of Charity	97
Charity, Self-Trust, and Trust in Others	102
BIBLIOGRAPHY	109

ACKNOWLEDGMENTS

I gratefully acknowledge Trent Dougherty's invaluable assistance as the director of this project, as well as the service of Michael Beaty, Alex Pruss, Francis Beckwith, and Jonathan Tran on the committee. I also thank friends and family too numerous to mention for their love, encouragement, and support.

PREFACE

Sometimes people fail to believe as they ought to believe, just as they sometimes fail to act as they ought to act. When they fail to act as they ought to, we often blame them for it. We do the same in the case of failures with respect to belief. We blame people for believing inappropriately. Let us call such blame *epistemic blame*.

My purpose in what follows is to gain an understanding of the nature and limits of epistemic blame. My central concern throughout is with epistemic blame as a social practice. In other words, my focus is not on what makes someone a fit target of epistemic blame, although plenty will need to be said about that. I am instead concerned with our situation when we blame others epistemically -- with what exactly we are doing in such situations, and how we can do it well.

I choose to focus on epistemic blame as a social practice because it is an important one. Our practices of epistemic blame help govern our relationships on an interpersonal and political level. If a friend unknowingly behaves towards me in a way that I find hurtful, how that affects my relationship with him may depend in part on whether I blame him epistemically for not knowing that I would find his behavior hurtful. If I have some serious ideological disagreement with someone, whether I blame her epistemically for believing as she does may determine whether I accept her as a reasonable interlocutor or ignore her.

On a larger social scale, judgments about what one should believe in a particular situation, and thus about what one can be epistemically blamed for not believing, are often key for determining legal liability. And whether we blame people epistemically for holding some political viewpoint can determine whether we grant them a seat at the table of political deliberation and compromise or dismiss them as cranks whose viewpoints we need not take into account.

These functions of epistemic blame make clear the importance of blaming appropriately. If we blame people epistemically when we should not, we run the risk of doing wrong -- harming relationships that should not be harmed, punishing people who should not be punished, and ignoring people who should not be ignored.

This investigation of the practice of epistemic blame will highlight the difficulty of doing it well. That difficulty, along with the harms involved in doing it badly, suggests the need for greater caution with our practices of epistemic blame than we often exhibit.

I will begin in chapter one by searching for an account of blame in general. This search is motivated by the fact that epistemic blame is a species of the genus *blame* -- epistemic blame is just that sort of blame that we direct towards people considered as epistemic agents. So we should have an account of blame in general within which to situate our account of epistemic blame, both in order to guide the account of epistemic blame and to allow for greater theoretical unity in our thinking about blame. In this chapter I will consider accounts of blame that appear in the literature and settle on a purely cognitive one. To blame people, I say, is simply to make a particular sort of judgment about them.

I continue in chapter two with the account of specifically epistemic blame. Building from the work in chapter one, I say that to blame someone epistemically is to judge that they have culpably failed to satisfy an epistemic norm. Fleshing out that account requires getting clear on what the epistemic norms are. I understand norms in relations to goals -- the norms governing an activity are what will help us accomplish the goal of the activity. So the epistemic norms are governed by the epistemic goal, which I hold to be getting at the truth and avoiding error. With that in mind I identify two kinds of epistemic norms. First there are what I call the responsibilist norms, according to which we should be conscientious and competent in our gathering of evidence. Then there is what I call the evidentialist norm, according to which we should believe in accordance with the evidence that we have.

In chapter three I ask whether epistemic blame is a kind of moral blame. Given how we have defined blame, that question amounts to asking whether epistemic norms are a kind of moral norm. I consider and reject some arguments inspired by John Locke and W. K. Clifford for the conclusion that they are. I then criticize some arguments in the contemporary literature for the claim that epistemic norms are not moral norms, and close by giving some reasons to think that they are.

Chapter four argues that we are not very good at epistemically blaming people, and often do it when we should not. This is because our cognitive access to the factors that determine whether someone ought to be blamed epistemically is rather limited. So we should be more cautious with epistemic blame than we tend to be.

In chapter five I try to articulate and defend a principle of charity to guide our practices of epistemic blame. We should, I say, assume that others are blameless if

possible. I defend this claim with arguments that draw on work by Quine on charitable interpretation and Foley on self-trust.

CHAPTER ONE

An Account of Blame in General

Introduction

Epistemic blame is a species of a genus – the genus *blame*. There are other species in that genus, most prominently moral blame. It would be to our advantage, then, to begin an account of epistemic blame by giving an account of blame more generally and then describing epistemic blame in terms of its specific difference. That is not to say that the only sensible way to talk about a species is in terms of its genus and specific difference. But to ignore blame more generally would risk leaving us with an account of blame that has no obvious connections to a concept to which it intuitively ought to have obvious connections. And the genus/specific difference approach makes for theoretical unity. It would be good to be able to talk about the various species of blame in a unified way, and that is best accomplished by talking about them in terms of their common genus.

Here is the account of blame I will be defending:

*X Z-blames Y for Φ -ing iff X judges that, in Φ -ing, Y culpably failed to satisfy some Z norm.*¹

¹ It is not within the scope of this project to give a formal account of the culpability condition. I take the account of blame to be largely neutral towards competing accounts of culpability. Some sort of "could have done otherwise" condition, very broadly construed, is probably the key element of most people's judgments about culpability. But for our purposes the main concern is that the account of culpability not itself include any talk of blame, on pain of circularity. I think that concern can be easily satisfied. "She knew it was wrong, and she didn't have to do it, but she chose to do it anyway" surely amounts to an ascription of culpability and makes no mention of blame.

There are two salient features of this account that are worth emphasizing. The first is that it is an explicitly *general* account of blame. It tells us what it is to blame someone morally, epistemically, or in any other way. To blame someone morally is to judge that they have culpably failed to satisfy some moral norm. To blame them epistemically is to judge that they have culpably failed to satisfy some epistemic norm. And so on.

The second salient feature worth emphasizing is that this is a purely cognitive account. It holds that a particular kind of judgment is both a necessary and sufficient condition for blame. Purely cognitive accounts are at odds with others in the literature. Other accounts include purely affective accounts of blame, on which to blame *X* is simply a matter of having a certain kind of affective reaction to *X*'s actions, and hybrid accounts on which blame involves both a judgment and an affective response. In what follows I'll describe and criticize some of these alternative accounts of blame. Then I'll respond to some objections to cognitive accounts and provide some positive arguments for them.

Competing Accounts of Blame

Blame as an Overt Act

One natural way to think of blame is as an overt action directed against the one blamed. On this view, the things we do to people we judge blameworthy—verbal expressions of disapproval, cold glares, withdrawal from social contact, etc—just are blame. You might think that this is so because to praise someone is to exhibit some sort of overt positive behavior towards them, and blame is often seen as just the negatively-

charged counterpart to praise.² If you have utilitarian inclinations, you might also see in this account of blame a nice utilitarian justification for the practice of blaming. People find being scolded and shunned unpleasant and might be encouraged to modify their behavior as a result of it. So if blame is just scolding, shunning, and the like, it's easy to see the utility of it.³

While this view of blame has at least one contemporary defender,⁴ it faces insurmountable difficulties. One such difficulty is that it seems that we can pretend to blame others. I could scold you and give you the cold shoulder because I disapprove of what you've done, but I could also scold you and give you the cold shoulder because I approve of what you've done and your hatred of me means that my outward signs of disapproval will encourage you to continue to act in that manner. In that case I would not be blaming you, which means that outward signs of disapproval are not sufficient for blame.

The obvious response is to modify the account to claim that blame involves outward signs of disapproval plus the belief that the target of those outward signs has acted badly. But again, I could have that belief and display those signs of disapproval without blaming.⁵ I could believe that you've acted badly, desire that you continue acting badly so that you will increase God's wrath against you, and verbally express disapproval

² George Sher, *In Praise of Blame*. (Oxford: OUP, 2005)

³ J. J. C. Smart, "An Outline of a System of Utilitarian Ethics," in *Utilitarianism: For and Against*, eds J.J.C. Smart and Bernard Williams (Cambridge, Cambridge University Press, 1973), 49-50, and P. H. Nowell-Smith, *Ethics* (London, Penguin, 1954), 306.

⁴ Richard Armeson, "The Smart Theory of Moral Responsibility and Desert," in *Desert and Justice*, ed. Sarena Olsaretti (Oxford, Clarendon Press, 2007)

⁵ Unless believing that you acted badly is itself sufficient for blaming you. But in that case, obviously, we don't need the overt action component anyway.

because I know that your hatred of me means that that will harden you in your ways. Doing so would not be blame. So outward signs of disapproval and the belief that the target of those outward signs has acted badly are not jointly sufficient for blame.

But a more fundamental objection to this approach to blame is that overt actions are not even *necessary* for blame. As Sher points out, we blame without showing any outward sign of the fact all the time. This is clearly true in cases in which we blame people who are spatially or temporally distant from us such that we cannot communicate with them at all. It is also true in the case of spouses, children, parents, and friends, from whom we may try very hard to hide our disapproval of something that they have done. It would be very strange to say that we blame them only insofar as we fail at that.

Affective Accounts of Blame

Blame is emotionally charged. Very often, when we blame someone, we are upset with them. We might take this as a reason to consider accounts on which such negative emotional reactions are constitutive features of blame. Let us call such accounts *affective accounts* of blame. This subsection will consider and criticize such accounts.

Blame as anger. The easiest way of constructing an affective account of blame would be to simply identify blame with an emotion. Anger is, of course, the most plausible candidate. After all, blame and anger seem to accompany one another, and they seem to be caused by the same stimuli. It would serve the virtue of theoretical simplicity if we simply dropped the idea that blame is a phenomenon separate from anger.

But that won't do, for a number of reasons. First, as Malle, Guglielmo, and Monroe point out,⁶ the language of blame and anger do not match up. We regularly speak of people as feeling anger, but we would never say that someone "feels blame." That suggests that we do not think of blame simply as an emotion. Moreover, blame requires warrant in a way that anger does not.⁷ We may judge anger to be irrational or vicious in the absence of justifying reasons, but we understand that it can exist without such reasons. "I don't know, I just am" is a perfectly coherent answer to the question "Why are you angry with her?" But as an answer to the question "Why do you blame her?" the response "I don't know, I just do" seems simply confused. Blame that is not for some identifiable wrongdoing or alleged wrongdoing is not merely irrational or vicious blame, it is not blame at all.

These considerations are enough to dismiss the idea that blame is merely anger.

Strawson. More sophisticated affective accounts are available. The *locus classicus* for such accounts of blame is P. F. Strawson's 1962 essay "Freedom and Resentment." Strawson describes blame in terms of negative "reactive attitudes," which are a particular kind of affective response to the actions of others. These reactive attitudes are embedded in and derive their importance from our social relationships with others:

The central commonplace that I want to insist on is the very great importance that we attach to the attitudes and intentions towards us of other human beings, and the great extent to which our personal feelings and reactions depend upon, or involve, our beliefs about these attitudes and intentions. [...] In general, we demand some degree of goodwill or regard on the part of those who stand in these relationships to us, though the forms we require it to take vary widely in different

⁶ Bertram F Malle, Steve Guglielmo, and Andrew E. Monroe, "A Theory of Blame," in *Psychological Inquiry* vol. 25, issue 2, (2014): 147-186.

⁷ *ibid*

connections. The range and intensity of our reactive attitudes towards goodwill, its absence or its opposite vary no less widely.⁸

It is important to us that the others whom we have to deal with be on our side. When their actions suggest that they aren't, we withdraw our goodwill from them. This withdrawal of goodwill is blame.⁹

There are a number of problems with describing blame as a negative affective response inspired by perceived lack of goodwill. For starters, it cannot easily be extended to give an account of what it is to blame others for something other than actions, such as their beliefs. I can blame others for their beliefs, and I may even resent them for their beliefs, but I can do these things without understanding their beliefs to involve either malice or apathy towards me. So if we want an account of blame in general that gives us insight into epistemic blame, we need to look elsewhere.

It also just is not the case that even blame motivated by actions always involves a perceived lack of goodwill where goodwill is important to me. Suppose that one of my devoted admirers were to murder someone who had failed to sufficiently appreciate my delightful whistling. I would, at least on my good days, blame my admirer. But her behavior clearly does not demonstrate any lack on her part of goodwill towards me. Likewise, none of Julius Caesar's actions demonstrate either hostility or apathy towards me, and even if they somehow had I can't imagine why I would care. Yet I can blame him for his treatment of the Gauls. Furthermore, it looks like I can blame people without any negative reactive attitude at all, regardless of whether that attitude is or is not

⁸ Peter Frederick Strawson, *Freedom and Resentment and Other Essays*. Routledge, 2008. Pp 5-7

⁹ It's worth noting that even this account involves *some* cognitive component. I have to believe that the person I'm blaming did such-and-such, and in doing such-and-such manifested a lack of goodwill towards me.

motivated by a perceived lack of good will. As Coates and Tognazinni (2012) point out, we can blame long-dead historical villains for their various crimes, but in many cases this blame is not accompanied by any particular affective response.¹⁰

There is one more argument against affective accounts that I will mention. Sher argues that affective accounts have no plausible way to account for blameworthiness.¹¹ To be blameworthy, on an affective account, would be to be such that it would be appropriate to have a negative emotional response towards you. But a negative emotional response from who? Certainly not just the person that you've wronged – we quite appropriately blame people who have not harmed us directly and have no personal connection with us, simply because we know that they are wrongdoers. On an affective account, that means that it is appropriate for us to be angry at, withdraw goodwill from, etc. anyone we know to be a wrongdoer. But all have sinned. We are all wrongdoers, and we all know that we are all wrongdoers. So, on the understanding of blameworthiness that an affective account of blame gives us, it is appropriate for everyone to be angry at and withdraw goodwill from everyone. But it is not appropriate for everyone to be angry at and withdraw goodwill from everyone. That means that the affective account gets blameworthiness wrong.

A proponent of affective accounts might respond by observing that the negative affective responses involved in blame needn't be particularly severe. So affective accounts are not committed to the claim that we must all walk around in a state of constant fury towards one another. Once we realize that the negative affective responses

¹⁰ Justin D. Coates and Neal A. Tognazzini. "The Nature and Ethics of Blame." *Philosophy Compass* 7, no 3 (2012): 197-207.

¹¹ Sher (2005) pp 85-88

involved can be of a very low-level sort, the conclusion that everyone appropriately has them towards everyone may seem at least a little less counterintuitive.

But it is still pretty counterintuitive. After all, affection and goodwill, not simply rather subdued animosity, seem to be what we ought to be aiming in our dealings with at least most people.

Tognazzini thinks that this criticism fails because it neglects the importance of standing when it comes to the appropriateness of blame.¹² The notion of blameworthiness that we get from the affective account of blame does not mean that everyone ought to always have negative reactive attitudes towards anyone else, because the mere fact that someone is blameworthy does not, by itself, mean that someone else has proper standing to actually blame them. We can lack standing to blame blameworthy others for a number of reasons. Perhaps doing so would be hypocritical because we ourselves are blameworthy in the same way. Or perhaps the matter is simply so distant from us (or so close to someone else) as to simply be none of our business. Responding to others who are blameworthy with negative reactive attitudes would, in such cases, be inappropriate. These and other such considerations of standing greatly reduce the scope of appropriate negative reactive attitudes. So the affective account's notion of blameworthiness does not leave us stuck with universal misanthropy after all.

It may well be the case that appealing to standing can help us ward off universal misanthropy. But I do not think that the picture here is as rosy as Tognazzini suggests. The appeals that he makes to standing, conjoined with the basics of Strawson's affective account, get the proper scope of blame wrong.

¹² Neal A. Tognazzini, "Blameworthiness and the Affective Account of Blame." *Philosophia* 41, no 4 (2013): 1299-1312.

Take lack of standing due to hypocrisy. Let us take it for granted that it would be inappropriate for me to blame others for wrongdoing if I myself am blameworthy in the same way. If Strawson is right, what does my blameworthy wrongdoing, in any case, ultimately consist in? In showing ill will to others, in failing to treat them as morally considerable. But why should the details of how I have done this be important? Smith has shown me ill will by stealing my bicycle. Why should the fact that I have not shown others ill will by stealing from them make it appropriate for me to have a negative reactive attitude towards Smith if I have shown others ill will in other equally significant ways? To hang the appropriateness of my negative reactive attitude on that is to rely on a distinction without a relevant difference.

So lack of standing to blame due to hypocrisy casts an implausibly wide net. I lack standing to blame others for wrongdoing if I am guilty of wrongdoing that shows a comparable degree of ill will towards others. That is not a plausible result. Negative reactive attitudes may indeed be inappropriate in such cases. But blame clearly is not. So this strategy of reducing the scope of blame by appealing to standing is not going to work.

So there are some strong objections to describing blame in terms of negative emotional responses. We will now move on to other possibilities.

Conative Accounts

Our range of possible psychological states is not exhausted by the categories of cognition and emotion, so dismissing affective accounts of blame does not leave us with no option but a purely cognitive account. We could try accounts of blame that make use

of psychological states like desires and intentions. Coates and Tognazinni call such accounts “conative accounts.”¹³

One such account is offered by George Sher. According to Sher, blame consists of a belief paired with a desire – the belief that another has acted wrongly and the desire that they not have done so.¹⁴ The desire component, he thinks, helps account for the motivational and emotional centrality of blame in our lives. After all the relevant desire, being a desire that the past be other than it is, cannot but be frustrated, and frustrated desires motivate us in all kinds of ways and provoke all kinds of emotional responses in us. And insofar as we regret being responsible for the frustration of the desires of others, this account explains why being blamed by others is so often an occasion for regret.

But, as Coates and Tognazinni point out, I can blame you for something and still not wish that you had not done it.¹⁵ If you get caught embezzling from the company, I could blame you for doing so and still be glad that you did it, in that having you out of the way greatly improves my own chances at getting that big promotion. If I learned that my parents had gotten together only after my mother had left some other man who was continually inconsiderate of her, I could blame that other man for being continually inconsiderate of her while still declining to wish that something that led to my own existence hadn’t happened. So it does not look as though the desire that the person blamed not have acted as they did is necessary for blame.

¹³ Coates and Tognazinni (2012)

¹⁴ Sher (2005) pg 12

¹⁵ Coates and Tognazzini. (2012) pg 201

Likewise, I can think that you've acted badly and wish that you hadn't, but for reasons that aren't of the proper sort to constitute blame. If you are my regular racquetball partner and you get arrested for bank robbery, I may well acknowledge that you acted badly in robbing the bank and regret that you have done so. But I may not care at all about the badness of your actions, and wish that you hadn't acted as you did only because finding a new racquetball partner is a hassle. If the belief that you acted badly in robbing the bank is not itself sufficient to constitute blame, adding my desire that you hadn't done so is not sufficient either.

So the desire that the person blamed not have acted as they did is neither necessary nor sufficient for blame.

T. M. Scanlon offers another conative account. On this account, "to *blame* a person is to judge him or her to be blameworthy and to take your relationship with him or her to be modified in a way that this judgment of impaired relations holds to be appropriate."¹⁶ The modification of the relationship here involves the modification of our intentions towards the blamed party, of our expectations about the blamed party, etc. Perhaps I originally intended to spend regular time with you, be kind to you, and refrain from violence towards you. But after you wrong me, I may judge that your wrongdoing has changed our relationship in such a way that those intentions are no longer appropriate, and alter them accordingly. Likewise, I may alter my expectations regarding the standard of behavior I can expect from you in the future, and thus withhold trust from you to some degree.

¹⁶ T. M. Scanlon, *Moral Dimensions: Permissibility, Meaning, and Blame*. (Cambridge, MA: Harvard University Press, 2008), 128-129

Of course, in the large majority of actual instances in which we blame people, we will in fact alter our intentions and expectations regarding them in some way or another. But that is just because if we believe that someone has wronged us, it is very often appropriate to do so. If you learn that I have wronged you, you gain information about my character, and it would only make sense to take that new information into account when you form expectations about my future behavior on the basis of which you decide how much to trust me. It would also make perfect sense to alter your behavior towards me, and thus your intended behavior towards me, in various ways. So as long as blame involves a judgment of wrongdoing, changes in intention and level of trust are going to follow naturally. We do not need to include them in an account of blame in order to account for their reliable correlation with blame.

And it seems pretty clear that we can blame people with whom we have no relationship to be altered – long-dead historical figures and contemporary strangers towards whom we have no particular intentions and no need to calibrate our level of trust. This is also true of people with whom we have had relationships but who are deceased. Suppose that, during one of our frequent hiking trips, I incompetently try to shove you off a cliff, lose my balance, and fall to my death. Your trust in me or lack thereof is no longer of much relevance, and the range of intentions you could have towards me is fairly limited. Perhaps you had previously intended to pray for the repose of my soul upon my death but decide that, given the circumstances, it would be appropriate to alter that intention. But surely you can blame me for my actions even if you manage to resist that temptation. So it does not look as though these alterations in intentions, expectations, etc. are a necessary feature of blame.

The Emotional Centrality of Blame

One of the most pressing worries about cognitive accounts stems from the fact that we know from experience that blame and negative affective responses are closely associated with one another. When we blame someone, we usually feel some degree of anger, dislike, resentment, revulsion, or so on towards them. George Sher expresses this worry:

‘[T]he belief account [...] utterly fails to capture the role that blame actually plays in our emotional lives. That blame plays some significant emotional role can hardly be denied. This is evident from the energy we expend in trying to affix it, from the rancor that often accompanies it, and from the urgency with which we seek to avoid it. [...] Thus, if the view that blame is a belief about a person’s moral balance or the state of his soul is to be at all tenable, its proponents must somehow explain why this is so.’¹⁷

If blame is a matter of belief, why is it always so deeply tied up with emotional responses?¹⁸ Of course, the fact that blame is deeply tied up with emotional responses isn’t a compelling argument that emotional responses are a component of blame. Things can be reliably correlated with other things without being nearly so conceptually related to them as that. But still, affective accounts can obviously explain the correlation quite nicely. If cognitive accounts can’t, then that puts them at a disadvantage relative to accounts that can.

An obvious response is to point out that the correlation between blame and negative affective responses isn’t perfect – take the case of emotionlessly blaming long-

¹⁷ Sher (2005). Pg 77.

¹⁸ Note that the beliefs involved in the cognitive accounts that Sher refers to are importantly different than the beliefs involved in the account that I’m defending.

dead historical figures, for instance. But the opponent of cognitive accounts may simply deny that these are genuine instances of blame. Such a denial would not be entirely *ad hoc*. After all, if the large majority of members of a class have a certain property, that gives us at least some reason to be skeptical of alleged members of that class that lack that property, unless we already have nailed down an account of the class that tells us otherwise. So it would be nice to have an explanation of why blame and negative affective responses are so closely correlated if they aren't conceptually related. That would at least reduce our skepticism about alleged instances of blame that involve no such affective responses.

Such an explanation is not difficult to come by. Recall the cognitive account introduced earlier – that X Z -blames Y for Φ -ing iff X judges that, in Φ -ing, Y culpably failed to satisfy some Z -norm. It is easy to see why judgments of that sort should fairly reliably be deeply tied up with negative emotional responses. Norms are, at least very often, going to be the sort of things we want to see satisfied. The satisfaction of norms is, after all, good, and we like goodness and want there to be lots of it. So when we judge someone to have culpably failed to satisfy some norm, we judge them to have culpably brought about the frustration of some of our desires. Anger, resentment, etc. are fairly natural responses in such situations. But that's a contingent psychological fact. It doesn't follow automatically from the judgment itself.

Likewise, many norms enjoin us to refrain from harming others in various ways. When we're in a position to notice violations of those norms, it will very often be because we or others that we care about have been harmed in some way. So judging that someone has violated those norms would mean judging that they have harmed us or

someone we care about. Again, anger and resentment will be common in such situations, but are contingent psychological responses rather than necessary features of the situation.

So it doesn't look like the frequent correlation of blame and negative affective responses gives us a strong reason to say that the affective responses are a component of blame.

Blame vs Blameworthiness

Here's another worry. You might think that the account I've given looks like a pretty good account of judging someone to be blameworthy, rather than of blaming. People are blameworthy if they culpably violate a norm, and if you judge them to have done so, well, you've judged them to be blameworthy. But judging someone to be blameworthy and actually blaming them are different things, so what's a good account of one can't be a good account of the other. So blaming must involve something else.

I agree that the account of blaming that I've given sounds like a good account of judging someone to be blameworthy. The question, then, is why we should think that judging blameworthy and blaming are distinct phenomena. If blame were an affective response they obviously would be, since judging someone worthy of an affective response and actually having that response are clearly different. But that would be to presuppose an answer to the question at hand. What other reasons could we have for thinking the two distinct?

The grammatical argument. One possibility is what we might call a grammatical argument. In general, that argument goes, judging someone to be worthy of being Φ -d is distinct from causing them to be Φ -d. That's just how judgments of worthiness work.

But it's not clear that we should accept that principle, despite its intuitive plausibility. Here's parallel principle that's also intuitively plausible but false: Saying that X is Y is different from causing X to be Y. To see that this principle is false, think of performative utterances. Under the right circumstances, to say that X is married or fired or excommunicated is just to cause X to be married or fired or excommunicated.

And the principle invoked in the grammatical argument seems to have a clear counterexample of its own. Judging someone to be worthy of being judged blameworthy certainly seems to be the same as causing them to be judged blameworthy. If that's right, then it's not true in general that judging someone to be worthy of being -d is distinct from causing them to be -d. So the grammatical argument doesn't seem very promising.

Partners in crime. Another option would be to describe a case in which it appears that someone is judged blameworthy but is not blamed. Coates and Tognazinni (2012) provide such a case. If I'm part of some joint criminal enterprise, my partners in crime might very well judge that I had violated some norm in stealing Smith's diamonds, and that I had done so culpably, but not blame me for doing so. They are, after all, my partners in crime. They want me to steal Smith's diamonds, and they praise me rather than blame me when I do so. So, judging someone blameworthy and blaming them are distinct.

There are a few possible responses to this case. One is to agree that it does show a difference between judging blameworthy and blaming, and alter the account in the following way:

*X Z-blames Y for Φ -ing iff X judges that, in Φ -ing, Y culpably failed to satisfy some Z-norm that X endorses*¹⁹.

“Endorse” is open to a few interpretations. Perhaps you endorse a norm if you consider it to be personally important to you, or if it’s a norm that you hold others to or desire or insist that others follow. Regardless, my partners in crime in the above example do not endorse the norm that I violated in stealing Smith’s diamonds, and that’s why they judge me blameworthy without blaming me.²⁰

That solves the problem, and all the work that I intend to do with the account of blame in general in the following chapters could be done just as easily using this modified account. But I’m not convinced that the alteration is necessary, because I think there’s a plausible argument to be had that the judgment that my partners in crime make about me is not the sort of judgment that my account describes as blame. I also think there’s a plausible argument to be had that if the judgment they make about me does count as blame under my account, then they do in fact blame me.

First, here’s the argument that my partners in crime do not make the sort of judgment that counts as blame under my account. Recall that that account crucially involves culpability – I must be judged not only to have violated the norm, but to have done so culpably. So the judgment that I have non-culpably violated a norm is not blame.

Now, one way to be non-culpable in violating a norm is to have overriding reasons to do so. Suppose I break the posted speed limit because I’m rushing a badly injured person to the hospital. I’ve violated the norm that one ought to obey traffic

¹⁹ I owe this suggestion to Alex Pruss.

²⁰ You may worry that this revised account is no longer purely cognitive. Perhaps it isn’t, and I don’t see much reason to insist on purity in that regard. But it’s worth noting that all the endorsement condition does is restrict the class of norms judgments about which count as blame.

regulations, but I'm not culpable because the reason I had for speeding – saving someone's life – outweighs my reasons for obeying that norm in this particular instance.

Judging that someone has culpably violated a norm, then, involves judging that they had no overriding reasons for acting the way they did. Such judgments will depend in part on the weight that the one making the judgment assigns to the norm that is being violated. Some people while acknowledging that there exists a norm that says that I ought to obey traffic regulations, will weigh the norm very lightly. So they might judge that I'm non-culpable in violating the norm if I'm running a bit late, or if I want to impress my passengers with my devil-may-care attitude, or whatever.²¹ Conversely, we could imagine someone so absurdly legalistic as to judge that even saving someone's life is not a significant enough reason to justify me in breaking the all-important norm that one should obey traffic regulations.

Apply this to my partners in crime. We've stipulated that they praise me for stealing Smith's diamonds. That would seem to indicate that they consider my actions praiseworthy. It's hard to see how they would do that without considering that I had some overriding reason to violate the norm that says I ought not steal other people's personal property. Such reasons might involve my own personal gain, the thrill of the crime, the well-being of the gang, or what have you. They're a gang of thieves, so it wouldn't be too surprising if they weighed personal property norms lightly enough that they would be overridden by any or all of those reasons. If that's right, then they judge that I violated the personal property norm, but they don't judge that I did so culpably. So their judgment

²¹ They'd be mistaken, but that's beside the point.

isn't one that counts as blame under the account I'm putting forward, and thus isn't a counterexample to that account.

But perhaps you don't think that they have to consider my actions to be praiseworthy in order to praise me. Perhaps you think it's entirely consistent to suppose that they might praise me while judging me to have culpably violated a norm. Even so, it's still not clear why, in that case, you should think that they aren't blaming me. We've stipulated that they praise me, but that's a poor reason to think that they don't blame me.

After all, the same action can satisfy some norms and fail to satisfy others. We can make judgments about both the satisfaction of one norm and the failure to satisfy the other. My stealing Smith's diamonds fails to satisfy the personal property norm, but it satisfies some thievery norms that my partners in crime recognize – perhaps norms that say that an act of thievery should be bold, or skillful, or remunerative, or stylish, or whatever. If they praise me it's because I've done well according to some norm that they recognize. That's totally consistent with them also judging that I've culpably failed to satisfy the personal property norm, and thus blaming me according to that norm. The praise rather than the blame seems to be what's governing their overall attitude towards me, but that's fine. It just indicates that they take the thievery norms to be more important than the personal property norms.

So, either the case of my partners in crime praising me for my misdeeds doesn't satisfy the conditions in the account of blame I've put forward, in which case it isn't a counterexample, or it does satisfy those conditions but there's no reason to think it's not an instance of blame. Either way, the case fails to undermine the account I'm defending.

The Problem of Forgiveness

Another potential problem for the account stems from the nature of forgiveness. Suppose that to forgive me for wronging you necessarily involves ceasing to blame me for the wrongdoing. On the account of blame that I'm defending, that would mean ceasing to judge that I had culpably violated a norm in wronging you.

But forgiveness doesn't seem to require that. Ceasing to judge that I had culpably violated a norm would mean either forgetting or excusing my actions, and forgiveness requires neither. Indeed, it seems reasonable to say that forgiveness is *incompatible* with either forgetting or excusing the thing forgiven. So, if forgiveness requires an end to blame, then there must be some other component of blame, either in addition to or in the place of the components mentioned in my account, the removal of which constitutes forgiveness.

There is some support in popular speech for the idea that blaming and forgiving are mutually incompatible activities. Psychologists speak of moving from blame to forgiveness, and we sometimes speak of blame and forgiveness as the two (presumably exclusive) responses we could make to being wronged.

But it's not clear that anything central to our notion of forgiveness is incompatible with blame. Forgiveness involves the refusal of the right to punish or demand recompense from the offender – think of a bank forgiving a debt.²² But there doesn't seem to be any reason to think that blaming necessarily involves insisting on such rights. Forgiveness involves the repairing of relationships ruptured or impaired by wrongdoing. So blame would be inconsistent with forgiveness if blaming necessarily

²² See Paul M. Hughes, "Forgiveness", *The Stanford Encyclopedia of Philosophy* (Spring 2015 Edition), Edward N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2015/entries/forgiveness/>>.

involved holding the relationship to be ruptured. But we've already dismissed accounts of blame that involve the rupturing of relationships. Forgiveness may involve ceasing to be angry at the wrongdoer, or at least disclaiming or distancing oneself from one's anger towards the wrongdoer. So forgiveness and blame would be incompatible if blaming involved having or endorsing a negative affective response towards the wrongdoer. But we've already rejected that account as well.

In general, it doesn't seem likely that there's going to be an "and no blaming" clause in whatever the correct account of forgiveness turns out to be. Whether or not forgiveness and blame are compatible, then, depends on what blame actually is. In other words, we would need an account of blame on hand that we could show was inconsistent with some element of our account of forgiveness. The claim that they are incompatible, then, looks question-begging as a premise in an argument against any particular account of blame. The argument from forgiveness, then, does not undermine the account of blame that I'm defending.

Degrees of Blame

Here's another worry. Blame comes in degrees. You can blame me a lot or a little. But the judgment that I've culpably failed to satisfy a norm does not come in degrees. You either make that judgment or you don't.

First of all, belief comes in degrees, so if a judgment is a kind of belief then the judgment that I have culpably failed to satisfy some norm can come in degrees. But that seems like the wrong response. It's not the judgment in blame that comes in degrees – the difference between blaming me a lot and blaming me a little is not in your degree of confidence in the blaming. It is rather in the degree to which you think I have

transgressed. I could fail to satisfy a norm by a lot or by only a little. I could fail to satisfy a norm of great importance or of minor importance. Degrees of blame can be accounted for perfectly well by judgments that I have culpably failed to satisfy a norm that involve more specific information about one or both of those factors. So the fact that blame comes in degrees does not present a problem for the cognitive account that I am defending.

Unreflective Blame

You might worry that the account of blame defended here presupposes far too much reflectiveness on the part of the blamer. Blame is often an instantaneous response to stimuli that takes place with no time to formulate any explicit judgment. We simply observe some bad behavior and begin to exhibit the various behaviors and emotional responses standardly associated with blame. Culpability and norms (never mind Z-norms) are somewhat abstract and rarified concepts, and surely people can blame even if they don't have those concepts available for deployment. In short, the cognitive account is just the sort of account a philosopher might be expected to give of a phenomenon that in fact takes place on a much more sub-rational level.

It is certainly true that we can blame someone without mentally formulating the sentence "That person has culpably failed to satisfy a norm." But some distinctions drawn from empirical psychology may be helpful here.

It is common in psychology to point out that while we can and do process information in a conscious, controlled way, we can also process it in ways that are unconscious and automatic. So, a novice driver may consciously think to himself "The car in front of me is stopping. I should hit the brakes. That's the pedal on the left. Gently,

though, don't hit it too hard." But an experienced driver will normally process and respond to this information without a moment's conscious thought, and is usually a safer driver for this fact. Only when conditions are such as to require a significant modification of normal driving behavior will an experienced driver consciously think through the mechanics of driving.

This processing of information is no less cognitive for being unconscious and uncontrolled. Likewise, we can say that the information relevant to judgments of blame is sometimes processed automatically and unconsciously. Thus Malle, Guglielmo, and Monroe, in their discussion of the judgments relevant to their own account of blame, say:

[T]here is no restriction built into [our model] regarding the modes of processing (e.g., automatic vs. controlled, conscious vs. unconscious) by which moral perceivers arrive at a blame judgment. Any given component's appraisal (e.g., about agentic causality or intentionality) may in principle be automatic or controlled, conscious or unconscious, depending on such factors as stimulus salience, existing knowledge structures, cognitive load, and so on.²³

With this in mind, we can see that this objection depends on a kind of over-intellectualization of judgment. The fact that we often blame instantaneously, without conscious inference or explicit mental formulation of the judgments involved, does not mean that blame is not a kind of judgment. It just means that we make judgments all the time without any conscious inference or explicit mental formulation of whatever we judge to be the case. We often make judgments of blame in an unreflective and automatic way, only subjecting those judgments to conscious scrutiny in difficult cases, or if we are asked to justify them.

But some may worry that this response is a bit too quick. After all, I have not merely said that blame involves some sort of cognitive processing -- the sort of claim that

²³ Malle, Guglielmo, and Monroe, 2014

would easily be satisfied if we blame in the same way that we respond to information about our environment while driving a car. I have claimed that blame is a particular kind of cognitive process -- a judgment involving norms and the culpable violation thereof. Why should we think that subconscious processing should count as that, even if it is real and genuinely cognitive?

It seems that for subconscious processing to count as a judgment about norms, there must be some real relationship between the norms and the processing. What sort of relationship could that be? Causal relationships seem like a promising possibility. An instance of sub-conscious processing can count as a judgment about norms if the subject's beliefs about the norms play a causal role in the sub-conscious processing

Such a causal role for beliefs about norms fit easily into the psychological literature on dual process theories. On such theories, sub-conscious processing is not purely instinctual and does not take place in a cognitive vacuum. Such processing is based on the pre-existing beliefs of the subject.²⁴ So, for instance, the quick and intuitive judgments people make about the validity of syllogisms are heavily influenced by their pre-existing beliefs about the truth or falsity of the conclusion. Subjects will quickly and intuitively judge a valid syllogism with an obviously false conclusion to be invalid and an invalid syllogism with an obviously true conclusion to be valid. More demanding, conscious cognitive processing is generally required to arrive at a correct judgment of such syllogisms. But for a valid syllogism with a clearly true conclusion or an invalid syllogism with a clearly false conclusion -- cases in which there is no conflict between the subject's pre-existing belief about the conclusion and the proper evaluation of the

²⁴ Wim De Neys, "Dual Processing in Reasoning Two Systems but One Reasoner," in *Psychological Science* 17, no. 5 (2006): 428-433.

argument -- subjects are able to arrive at the correct judgment about the syllogism without any such conscious processing.²⁵

A subject's beliefs, then, influence the outcome of sub-conscious cognitive processing. If beliefs about norms influence the outcome of the sub-conscious cognitive processes involved in unreflective blame, that looks like a substantial enough connection between the norm and the judgment to satisfy this account.

It is worth noting that, as in the example above, sub-conscious cognitive processing is often a rather inferior way of getting at the truth. That is not a problem for the view I have been defending -- nothing about the claim that blame is the judgment that someone has culpably violated a norm implies that all such judgments will be truth-apt. But it does give us a reason to be particularly cautious when we find ourselves blaming unreflectively.

Thus ends the consideration of possible objections to the account that I have put forward.

For the Account

Of course, nothing that has been said so far constitutes any kind of positive argument for the account of blame that I've been defending. What can be said in favor of the account?

Certainly, there is a good case to be made for including negative normative judgments in an account of blame. Judgments that someone has culpably failed to satisfy a norm are just fault-finding, and blame without fault-finding seems incoherent. The faulty light fixture in my kitchen is a frequent target of negative affective responses, but

²⁵ *Ibid*

has never yet been a target of blame. It can't be, precisely because it can't be at fault for anything. Similar negative affective responses directed at humans, in the absence of any judgment that they are at fault, look like they are in the same boat.

Besides, if negative normative judgments aren't going to be part of our account of blame, then it looks like we could blame even while explicitly rejecting such judgments – while explicitly holding that the person we blame is not blameworthy. “You are not blameworthy, but I blame you” sounds incoherent, though, or at least very strange.

That is a reason to make the judgment that someone has culpably violated a norm part of the account of blame, but it is not much a reason to make it the only part. The necessity of such judgments for blame is no argument for their sufficiency.

Other than the failure of the various proposed other requirements for blame, I think that the main reason for accepting the account is simply that it genuinely is an account of blame in general. It probably has not escaped your notice that the other accounts of blame discussed in this chapter are all very much accounts of moral blame. It is not easy to see how any of them could be extended to other kinds of blame. But it is quite easy to see how blame as the judgment that one has culpably failed to satisfy some norm could be modified to generate accounts of specific kinds of blame. Moral blame is the judgment that one has culpably failed to satisfy a moral norm. Epistemic blame is the judgment that one has culpably failed to satisfy an epistemic norm. There are as many kinds of blame as there are kinds of norms that one could be judged to have culpably failed to satisfy, all of them neatly tied together under the auspices of a straightforward account of blame in general. If we take seriously the idea that there are multiple kinds of

blame that are all equally blame, then the theoretical unity provided by this account is a significant advantage.

Of course, we will never encounter blame in general in the wild. All blame is blame of some specific kind or another. We can't test the account by seeing how well covers agreed-upon instances of blame in general, since there aren't any. If we're going to test anything against our intuitions about actual cases, it will have to be the accounts of particular kinds of blame. As nice as the theoretical unity provided by the account of blame in general is, it won't be worth much if the accounts of specific types of blame that it generates aren't plausible and fruitful. So we will have to go into the details of the account of epistemic blame and see if it works. It is to that task that we now turn.

CHAPTER TWO

What is Epistemic Blame?

Introduction

In this chapter I will discuss the nature of specifically epistemic blame. I need to preface this discussion with an important distinction. You might quite naturally expect me to connect epistemic blame to the notion of justification. After all, justification and blame seem like they ought to be closely connected. So you might expect that epistemic blameworthiness has something to do with a lack of epistemic justification, and that epistemic blaming is going to somehow involve judging that a particular belief lacks epistemic justification.

I will not go that route. I want to entirely avoid talk of epistemic justification. That's not to say that a blameless belief, on the account of epistemic blame that I will put forward, is not one to which the English word "justified" could reasonably be applied, nor that a blameworthy belief couldn't reasonably be called "unjustified." But justification, in modern epistemology, is very closely tied to the concept of knowledge. Justification is a normative component of knowledge. The need to have it play that role constrains much of the contemporary discussion of justification's nature.

There is no need for a discussion of the nature of epistemic blame to be subject to the same constraint. After all, one might think knowledge quite unimportant, or even dismiss the concept entirely, but still feel the need to talk about epistemic blame.

There is also no need for the discussion of epistemic blame to place any constraints upon our theory of knowledge. If the correct account of epistemic blame implies that a belief that fails to satisfy some theory of knowledge's normative component could nevertheless be blameless, it would be strange to draw from that fact the conclusion that that theory of knowledge is false. So in general, we need to keep the discussion of epistemic blame untangled from worries about knowledge. To accomplish that, we should leave justification to the knowledge theorists and go our own way without it.

From the previous chapter, the basic account of epistemic blame that I will defend should already be clear:

*X epistemically blames Y for Φ -ing iff X judges that, in Φ -ing, Y culpably failed to satisfy some epistemic norm.*¹

Fleshing out the account requires getting clear on what epistemic norms are.

Norms and Goals

What makes a norm epistemic? The norms of a given activity are governed by the goal of that activity – the norms are what they are because of their connection to achieving the goal. If I am cooking pasta, the goal of my activity is to produce some tasty, nutritious pasta. The norms to which that activity is subject are governed by that goal. I ought to bring the water to a vigorous boil before I add the pasta. I ought to add a little salt. I ought to refrain from cooking the pasta in rubbing alcohol or cement. Why is

¹ Note that on this account, while we would normally expect epistemic blame to attach either to beliefs or to some action involved in the belief-forming process, if it turns out that something that is neither a belief nor directly involved in belief formation can nevertheless violate an epistemic norm, it is a fit target for epistemic blame.

my activity subject to these norms? Because following them will help, and failing to follow them will hinder, me in reaching my goal in cooking pasta.

Of course, while making pasta, I am subject to all kinds of norms that are not governed by the goal of making tasty, nutritious pasta. I should not obtain the pasta by theft. I should breathe and refrain from blaspheming while cooking the pasta. I should wash the pot after I'm done. All of these are norms to which my pasta making is subject, but they are not governed by my goal in making pasta. That might seem to be a problem for the idea that the norms proper to an activity are governed by the goal of that activity.

But, while my pasta making is subject to all those norms, they are not properly pasta making norms. After all, even when I am making pasta, the goal of producing some tasty, nutritious pasta is not the only goal that I have. I also desire to remain alive, morally upright, tidy, etc. These norms are governed by those goals. They are thus moral norms, staying-alive norms, etc. That they impose certain constraints on my pasta making does not change that. So the fact that my undertaking of an activity can be subject to norms not governed by the goal of that activity does not undermine the account of the relationship between norms and goals given here.

Here is another worry about saying that the norms of an activity are those governed by the goal of that activity. Suppose I made the pasta from scratch, using only the most premium ingredients. Suppose I cooked it in water from a spring in rural Sicily that contains just the right balance of minerals for a perfect flavor and texture. Suppose I spend \$300 on the ingredients for the sauce. All of these things would, no doubt, help me attain the goal of producing some delicious pasta, and my pasta will be less delicious if I omit them. But surely I don't violate any pasta making norms by being less extravagant

than that. I am not properly subject to pasta blame for making a merely ordinary batch of pasta.

There are two ways to respond to this worry. One involves being more specific about the goal of the activity. Rather than saying that my goal is to produce delicious pasta, we could say that my goal is to produce pasta that clears some threshold of deliciousness. The norms of pasta making will only be what is required to get me to that threshold. Anything that would merely increase the distance by which I clear the threshold is not going to be a pasta making norm. What precisely the pasta making norms are, then, will depend on how good I want the pasta to be.

Another response is to agree that all of those extravagances constitute norms of pasta making, but add that I am not properly subject to pasta blame if I omit them if I am justified in doing so by the degree to which fulfilling the extravagant norms would interfere with other goals that I have. So even though there is some sense in which I *should* make the pasta from scratch, I am excused from doing so if investing the time, money, etc. that would be required to do that would interfere with other goals that I consider more important than making delicious pasta.

So, norms are governed by goals. The epistemic norms, then, are going to be those norms the observing of which will help us attain the epistemic goal.

Epistemic Norms and the Epistemic Goal

Figuring out what the epistemic norms are will therefore require us to say something about the epistemic goal. That is, of course, a controversial issue. There are quite a few proposed epistemic goals to be found in the literature. Perhaps the epistemic goal is to acquire true beliefs and avoid false ones, or to form beliefs via a process that is

objectively likely to produce true beliefs rather than false ones. Perhaps it is to believe justifiedly rather than unjustifiedly. Perhaps it is to acquire beliefs that count as knowledge and avoid beliefs that don't. Perhaps it is to increase reproductive success, or to acquire something like wisdom, understanding, or theoretical insight.

We can have all of these goals, and perhaps more besides, when we put our cognitive equipment to work.² Maybe all of these goals govern norms that, because they are governed by goals involving the use of our cognitive equipment, could reasonably be called epistemic. But looking at our actual practices of epistemic blame should give us a decent sense of which goals are relevant for our purposes.

We do not blame people epistemically for using their cognitive equipment in a way that is harmful to their reproductive fitness, and beliefs may be fit targets of blame even if they increase reproductive fitness. So reproductive fitness doesn't seem to be the goal that governs the norms under which we ascribe epistemic blame.

Demon world considerations should suffice to show that reliability isn't the goal that governs the norms under which we blame. In a demon world a belief-forming process may reliably produce false beliefs yet be blameless, or reliably produce true ones yet be quite blameworthy.

Perhaps we do sometimes blame people for failing to attain wisdom, understanding, insight, or something along those lines. Calling someone a fool certainly sounds like reproach, and it is reproach targeted at their failure to attain wisdom or understanding.

² See, for instance, the argument for pluralism about the epistemic goal in Jonathan Kvanvig, "Truth and the Epistemic Goal," in *Contemporary Debates in Epistemology* ed. Matthias Steup and Ernest Sosa (Malden, MA: Blackwell, 2005), 285-296

But the blame we are concerned with is blame that attaches, at least primarily, to discrete beliefs. Recall the moral legal purposes for which we need a notion of epistemic blame. Wisdom, understanding, insight, and the like are not a matter of discrete beliefs. So the goal of attaining wisdom, understanding, or insight does not seem to be the goal we are after.

That leaves three main options – the goal of attaining knowledge, the goal of believing justifiedly rather than unjustifiedly, and the goal of getting to the truth and avoiding error.

As for attaining knowledge, we do not blame people for being gettiered. Indeed, it seems impossible that there could be a don't-be-gettiered norm involved in our belief formation, since the whole point of Gettier cases is that the subject has done everything that could reasonably be asked of them but is robbed of knowledge by circumstances they have no access to. So, insofar as not being gettiered is a necessary condition for knowledge, it looks like we don't blame people for failing to acquire knowledge.

Believing justifiedly rather than unjustifiedly does seem like something that we want to accomplish when we operate our cognitive equipment. But in general, we want our beliefs to be justified because we want them to be true.³ It is the fact that, under normal circumstances, justified beliefs are more likely to be true than unjustified ones

³ See Marian David, "Truth as the Primary Epistemic Goal: A Working Hypothesis," in *Contemporary Debates in Epistemology* ed Matthias Steup and Ernest Sosa (Malden, MA: Blackwell, 2005), 296-312

that make them desirable. So it seems as though justified belief is a secondary rather than primary epistemic goal.⁴

That leaves the goal of getting to the truth and avoiding error as the goal that governs the norms under which we blame epistemically. There are multiple ways of specifying this goal. Does it mean getting to any truth at all, or only some set of important or interesting truths? How do we prioritize the two halves of the goal? I could be quite risk-averse in my believing and withhold belief in almost all circumstances. I would therefore greatly increase my chances of avoiding error, but at the cost of greatly decreasing my chances of getting to the truth. Or I could adopt a strategy of believing everything I hear except in cases in which two things I have heard flatly contradict each other. That would greatly increase my chances of getting to the truth, but at the cost of greatly reducing my chances of avoiding error. Or I could adopt some policy between those extremes. Which policy best represents our epistemic goal?

These are important questions in themselves, but they are not important questions for our present purpose. In any case the basic means we use to get to the truth and avoid error will be the same. Questions about how much truth we want to get to and how much error we want to avoid are questions about how much we ought to employ the means, not about what the means themselves are.

That the relevant epistemic goal is getting to the truth and avoiding error cannot be taken to mean, though, that people are going to be blameworthy every time they fail to get to the truth and avoid error (or that they won't be blameworthy any time they do so.) That is because in some cases the best way to get to the truth and avoid error is going to

⁴ But if you want to say that justified beliefs are the primary epistemic goal, then I can accept that claim without changing what I say below about the epistemic norms, since the recommended procedures for acquiring justified beliefs are identical to those for acquiring true beliefs.

be opaque to us – in a demon world, to take only the most extreme example – and thus we can't be considered culpable for failing to do the things that would help us get to the truth and avoid error.

The Epistemic Norms

Blaming someone for not doing what it takes to get to the truth and avoid error, then, involves making some judgments about what ways of getting to the truth and avoiding error they ought to be aware of. There are two important ways of getting to the truth and avoiding error that, in normal circumstances, sufficiently reflective people ought to be aware of. One is to gather evidence in a way that is conscientious and competent. The other is to believe in a way that is licensed by the evidence so gathered.

So, there are two broad classes of epistemic norm. One class contains the norms governing inquiry. I will call these the responsibilist norms. The other is simply that we believe in a way that fits our evidence. I'll call this the evidentialist norm. I will discuss each in turn.

The Evidentialist Norm

The main question that arises concerning the evidentialist norm is what, precisely, the evidence in accordance with we are supposed to believe is. There is a very important constraint on the notion of evidence involved in the evidentialist norm. Whatever evidence is, it must be the sort of thing our failure to believe in accordance with could be accessible to us. That is because epistemic blame involves not only the judgment that one has failed to satisfy some epistemic norm, but that one has *culpably* failed to satisfy it.

But culpability requires a *mens rea*, or a guilty mind.⁵ If I'm going to be culpable for doing something, I have to be aware (or at least be such that I ought to be aware) both of the wrongness of what I do and of the fact that I am doing it. So I can only (rightly) be blamed for failing to believe in accordance with the evidence if I should be aware that that is what I am doing.

In the fourth chapter I will argue that the any notion of evidence that meets this requirement will be something like what Earl Conee and Richard Feldman call "seeming evidentialism," according to which "someone's evidence about a proposition includes all that seems to the person to bear on the truth of the proposition."⁶

The Responsibilist Norms

The way that we go about gathering evidence is just as important as what we do with that evidence once we have it. It is entirely possible to believe in accordance with the evidence that you have but still be blameworthy because the evidence that you have is, through your own fault, not of high quality. So, there are norms governing the process of gathering evidence – in other words, norms of inquiry. I call these the responsibilist norms, since they enjoin responsible inquiry and are the sorts of norms that are of particular interest to responsibilists.

The most obvious and basic norm of inquiry is that inquiry should happen. We should not be content with whatever meager body of evidence we have managed to find ourselves with through no particular effort. Relevant sources should be consulted.

⁵ See Nikolaj Nottelmann, *Blameworthy Belief: A Study in Epistemic Deontology*. Vol. 338. Springer Science & Business Media, 2007.

⁶ Earl Conee, "First Things First," in *Evidentialism: Essays in Epistemology* ed. Earl Conee and Richard Feldman (Oxford: Clarendon, 2004), 11-36.

Relevant inferences should be performed. Relevant experiments should be undertaken. Bodies of evidence that result from such processes of inquiry are more likely to support true beliefs, and less likely to support false ones, than bodies of evidence that just sort of happen.

But quality of inquiry matters as well as quantity. We cannot simply spend X amount of time engaged in just any activity of inquiry and declare that all the responsibilist norms have been satisfied. I do not intend to provide a complete taxonomy of all the norms involved in high-quality inquiry, but some are obvious. We should consult reputable sources and avoid obvious quacks. We should make an effort to avoid only consulting sources that would affirm our pre-existing ideas or tell us that what we wish to be true is true. If we make any inferences or perform any experiments, we should do so according to the normal standards that govern inferences and experiments. We should do all of these things because, again, evidence that results from inquiry conducted in this way is going to be more likely to help us get to the truth and avoid error than inquiry conducted less conscientiously.

The evidentialist norm interacts with the responsibilist norms in important ways. How much inquiry I need to perform on a question depends on the quality and thoroughness of the evidence I have that bears upon that question. I already have great evidence that I have hands, and unless something happens to shake my confidence in that evidence I am under no particular obligation to inquire any further into the question. The same is true of visual beliefs in general. The evidence about the world that our eyes deliver to us with no particular effort on our part is, under normal circumstances, quite sufficient for the sorts of beliefs that we base on that evidence. No further inquiry is

needed, because the body of evidence is already as good as it needs to be. More inquiry would have a negligible effect on my likelihood of getting to the truth and avoiding error.

The same point extends to what we might call meta-evidence – evidence about the quality of a body of evidence that bears on a specific question. No matter what the actual quality of my body of evidence is, I might have meta-evidence that renders me blameless in deciding that my evidence is good enough to justify cutting off further inquiry. If so I shouldn't be blamed for cutting off inquiry, even if more inquiry would in fact significantly improve my chances of getting to the truth and avoiding error.

The evidence we have about our own processes of inquiry also affects what further inquiry we need to do. If I very vividly recall asking a very honest expert on American history when the Civil War ended and being told that it ended in 1864, I can't be blamed for closing inquiry into the question of when the Civil War ended even if this memory is a false one implanted in my head by an evil demon. If I have evidence that renders me blameless in thinking that a source is reliable, I can't be blamed for focusing my inquiry on it even if it is in fact a pack of vicious lies. The evidence that we have determines what sort of inquiry is demanded of us.

The Other Norms

Both the responsibilist and the evidentialist norms are closely related to the formation of beliefs. It thus feels perfectly comfortable to describe blame under those norms as epistemic blame. Both are, after all, clearly norms governing how we go about using our cognitive equipment. But there are all kinds of ways I could help myself get to the truth and avoid error that are much less directly connected to the belief-forming process. Drinking a whole fifth of bourbon in the space of an hour, gouging out my eyes,

and plenty else besides are all going to reduce my ability to get to the truth and avoid error. By doing such things, do I violate any properly epistemic norm? It is at least a little odd to say that I do, despite the fact that maintaining my sobriety and the proper functioning of my sense organs are important means of attaining the epistemic goal.

I know, for instance, that if I fail to wear my glasses, I will miss out on a lot of true beliefs about the external world and probably pick up some false ones that I would not have had I worn my glasses. It certainly seems as though I ought to wear my glasses. But it seems at least a little odd to say that there is an epistemic norm according to which I ought to wear my glasses, since putting on my glasses is not an operation of my cognitive equipment. The question, then, is whether *any* norm that is governed by the goal of getting to the truth and avoiding error counts as an epistemic norm, or only those having to do with the operation of our cognitive equipment.

One option would be to connect the norm more closely to belief-forming processes by observing that while not wearing my glasses is not an operation of my cognitive equipment, forming beliefs about the external world while not wearing my glasses is. Perhaps that is the relevant norm here. If it is, then we don't have to say that I do anything *epistemically* blameworthy by not wearing my glasses.

This approach has some advantages. It does seem clear that we might blame me in some circumstances for forming beliefs about the external world while not wearing my glasses. But why? Because in a lot of cases, if I did so I would be forming beliefs using faculties that I have good reason to distrust the deliverances of. That can be handled easily by the evidentialist norm, though. If I have evidence that the evidence produced by a particular faculty is likely to be misleading, that serves as an undercutting defeater that

greatly reduces the evidential force of the evidence delivered by that faculty. So, there is a simplicity advantage to dealing with the question in this way. If there is no epistemic norm that says that I ought to do all the things that aren't directly related to the operation of my cognitive equipment but that nevertheless help me get to the truth and avoid error, but rather norms that say that I ought not believe on the basis of certain faculties unless I have done so, then we can describe these norms as a particular version of the evidentialist norm.

But I don't think that we should go this route. After all, it looks like I can be subject to properly epistemic blame for not wearing my glasses even if I don't form any sort of beliefs about the external world. Suppose that you wave at me while passing about thirty feet away from me on the quad. I think that there's someone waving at me, and think that it might be you. But since I'm not wearing my glasses I know that my visual faculties are not very reliable, so I withhold judgment on the question and ignore you. You are insulted and ask me about it later. "Oh, I say, I wasn't wearing my glasses, so I thought it best to suspend judgment on whether or not that was someone waving at me." It seems like your natural response would be "Well, if your vision is so bad that you need glasses to discern whether or not someone thirty feet away is waving at you, then you ought to wear your glasses every time you leave the house." That blame does not turn on my forming beliefs, but on my being inexcusably unprepared to form appropriate beliefs.⁷

⁷ This makes sense, if we take seriously the view that the epistemic goal is to get to the truth and avoid error. Believing when I have good evidence that my evidence is unreliable is unwise given the last half of that goal, but not taking steps to ensure that my faculties are operating at a reasonably high level is unwise given the first.

So I think that we must accept that epistemic norms can include things that we should do or avoid in order to improve our chances of getting to the truth and avoid error that are not directly involved in the operation of our cognitive equipment.⁸

This result is not as surprising as it may seem. Proper use of a thing includes proper maintenance of it. Suppose that I consistently run red lights while driving. Doing so would violate some of the norms of driving. If I defended myself by claiming that I really was an excellent driver, and that I had no choice but to run the red lights since my brakes haven't worked for the past six months, that would be an odd sort of defense. A good driver, you would not doubt reply, keeps his or her car in a reasonable state of repair. I may still be a *skilled* driver, if I drove the brakeless car as well as you could reasonably expect anyone to drive a brakeless car, but I would not be a *good* driver. I would still not be a good driver even if every intersection I hit had a green light, and I thus never violated any norms related to running red lights or any other traffic rule. The norms of driving extend beyond the actual operation of the vehicle. Good drivers get their brakes fixed.

Likewise, I may get quite skilled at navigating the array of blobs that is my visual field when I'm not wearing my glasses. I may train myself to get the maximum amount of reliable information about the world that could be gotten from such visual experiences. But that would not make me a good operator of a visual apparatus. To be that, I would need to take the easy steps necessary to ensure that my visual apparatus actually works as intended.

⁸ As far as the taxonomy of norms go, I intend to treat these as responsibilist norms, since they concern things that we should do in order to ensure that we get high quality evidence on which to base our beliefs.

So I think that we should accept that there are epistemic norms that do not directly involve the operation of our cognitive equipment. The worry about that conclusion is that it seems to give epistemic norms a foothold in most, if not all, areas of life. Extreme cases of self-induced cognitive impairment like those discussed above do not, after all, exhaust the things we can do that do not directly involve the operation of our cognitive equipment but nevertheless affect our ability to get to the truth and avoid error.

Some diets are better for brain function than others. Do I violate an epistemic norm by not choosing the most brain-healthy diet that I am aware of? Staying in one location after I've already taken in all the sensory information that I'm likely to about it prevents me from learning new things about a less familiar location. Writing down philosophical ideas that I've already worked through rather than reading an encyclopedia prevents me from forming some new true beliefs. Falling in love with somebody would likely leave me less likely to recognize certain unflattering truths about them. Sleeping any more than is necessary to maintain good brain function simply wastes time that could be spent learning things. Risking my life to save another means intentionally risking a state of affairs – my death – that would seriously hamper my ability to get to the truth.⁹ (Although avoiding error would be a cinch.) If I fail to put on my glasses immediately upon waking up because I'm too lazy to go get them and I don't feel the need for them, I willingly choose to deprive myself of some important visual information for whatever space of time it takes me to go get them. If I refrain from inquiring into the sex lives of strangers, I miss out on the possibility of acquiring some valuable truths.

⁹ Assuming that there's no cognitively advantageous afterlife.

It seems absurd to say that someone would be epistemically blameworthy for all these things, but if we allow that there are epistemic norms not directly connected with the operation of our cognitive equipment, it's difficult to avoid the conclusion that doing these things can violate epistemic norms. And if doing those things can violate epistemic norms, then I could be epistemically blameworthy for doing them.

But here again, the fact that our account of epistemic blame requires that the violation of the norm be *culpable* is helpful. One is not culpable for violating a norm if they have sufficient, overriding reasons to do so. And that seems to be plausibly the case in the scenarios described above. The human cognitive apparatus may have as its primary goal the acquisition of truth and the avoidance of error, but that does not mean that that is the primary goal of the human. Some of our other goals may come into conflict with the goal of getting to the truth and avoiding error. When that happens, there is no reason to suppose that the epistemic goal should automatically win. In all of the cases given above, I refrain from maximizing my chances of getting to the truth and avoiding error for the sake of some other good. If these other goods are sufficiently weighty as to give me an overriding reason to violate the epistemic norm, then I do so inculpably and am not a fit target of epistemic blame.

That should help defuse the worry that the way we've described epistemic norms leaves a vast number of clearly unobjectionable activities epistemically blameworthy. But it brings with it worries of its own.

The first is simply that it might seem unsatisfying to say that choosing various everyday goods over doing absolutely everything possible to get to the truth and avoid error is a violation of an epistemic norm, but an inculpable one. To say that an agent has

violated a norm, but done so inculpably, is to say that the agent has done something genuinely bad but is for whatever reason excused. If you say that my extreme hunger renders my stealing food inculpable, you aren't saying that the theft was good or even neutral. You are simply saying that circumstances are such that I should not be blamed for doing a bad thing.

But in the cases described above, it seems odd to say that I am doing a bad thing, but am exempt from blame because of the overriding reasons for doing so. Refraining from inquiring into the sex lives of strangers is not bad but excusable. It is, under most circumstances, good or at worst neutral. If I avoid a diet that maximizes brain function because I think that a different diet is healthier in other ways, more ethical, or even simply more pleasant, that is at least possibly a good or neutral thing to do. It does not seem like the sort of thing that would need to be excused, in the way that non-culpable violations of norms are excused.

If you find this worry compelling, one way to defuse it would be to alter the description of the epistemic norms such that if there are overriding reasons not to do whatever it is that would maximize the chances of getting to the truth and avoiding error, the norm is not violated at all. Take brain health. Perhaps the relevant epistemic norm with respect to maintaining brain health is not that I should adopt whatever diet is likely to maximize my brain's ability to get to the truth and avoid error. Perhaps it is that I should adopt whatever diet is likely to maximize my brain's ability to get to the truth and avoid error *if there are no overriding reasons to do otherwise*. If that's true, the choosing a diet that is non-optimal for brain health if I have a good reason to do so is not a non-culpable violation of the norm. It is not a violation of the norm at all.

Well and good, but how does that fit with our description of the relationship between epistemic norms and the epistemic goal? If the norms of an activity are governed by its goal, what justifies the *if there are no overriding reasons to do otherwise* clause? That look like it's letting non-epistemic factors constrain norms that should, on the account we've given, be governed only by the epistemic goal. In other words, it is doubtful whether norms with an *if there are no overriding reasons to do otherwise clause* are, on our account of the relationship between norms and goals, properly epistemic.

So, let's return to the epistemic goal. If norms with *no overriding reasons* clauses are going to be properly epistemic, we need an analogous clause in the epistemic goal itself. The epistemic goal would need to be to get to the truth and avoid error if there are no overriding reasons to do otherwise. Could adding such a clause be anything but hopelessly *ad hoc*?

Yes, it could. Take our brains, or sensory organs, and anything else you might choose to consider part of our cognitive equipment. Why does that equipment have the function that it does? Why are my eyes for getting to the truth and avoiding error about the external world, instead of for enabling flight or being a manly shade of blue or what have you? One plausible theory is an etiological account that appeals to past natural selection. On that view, my eyes are for helping me learn things about the external world because that is what my very distant ancestors' eyes did that secured those ancestors a reproductive advantage.¹⁰

Well, what did my distant ancestor's brains and other pieces of cognitive equipment do to secure them a reproductive advantage? Help them get to the truth and

¹⁰ Ruth Garrett Millikan, "In Defense of Proper Functions," *Philosophy of Science* Vol. 56, No. 2 (1989): 288-302.

avoid error? Yes, but subject to certain constraints. Only so much getting to the truth and avoiding error actually contributes to an organism's reproductive success. And brains that devote resources to getting to the truth and avoiding error beyond the point at which it ceases to contribute to the organism's reproductive success are, from an evolutionary standpoint, wasting resources that could be used to boost the organism's reproductive prospects in other ways. An organism with such a brain would actually be at a disadvantage relative to one that devoted just enough resources to getting to the truth and avoiding error, but no more than that. So there's a plausible evolutionary story to tell on which the function of our cognitive equipment is to get to the truth and avoid error while appropriately balancing that task with other needs that the organism has.

A theistic account of the function of our cognitive equipment can make the same move. On a theistic account, the function of our cognitive equipment is owed to God's purpose in designing it, much as the function of any human-made artifact is owed to the purpose of its designer. But God, in designing our cognitive equipment, would not want it to pursue the task of getting to the truth and avoiding error in a way that compromised our other basic needs. So on a theistic account of biological function, it again looks as though the function of our cognitive equipment is going to be to get to the truth and avoid error *subject to constraints posed by other needs that we have*.

So it is not *ad hoc* to include a *no overriding reasons to do otherwise* clause in the description of the epistemic goal. That means that a descriptions of epistemic norms that include analogous clauses fit quite well with the account we have given of the relationship between norms and goals. That, in turn, means that we can appeal to epistemic norms described in that way in order to avoid having to say that failing to

maximize our chances of getting to the truth and avoiding error is a violation of an epistemic norm even if we have good overriding reasons to do so. That should deal with the first worry about appealing to overriding reasons to do otherwise in order to avoid concluding that a vast range of clearly unobjectionable activities are epistemically blameworthy.

The second worry is that such a solution goes too far and leaves open ways to declare blameless some things that are clearly epistemically blameworthy. We said that I am not epistemically blameworthy for refraining from inquiring into the sex lives of strangers because I had overriding reasons (social, legal, emotional, etc) for so refraining. That seems clearly right, but the same sort of reasoning can be put to work in more problematic cases.

Imagine Sally the agnostic, who is an agnostic primarily because she has never given much thought at all to arguments for or against God's existence. Suppose you give her a short statement of the major arguments for and against God's existence. She's quite capable of understanding and evaluating them, and she has no other pressing demands on her time. But she refuses to inquire into those arguments. She wishes to ensure that she remains an agnostic, because she believes that being an agnostic makes her seem more interesting at parties.

You call her out for her intellectual apathy. She defends herself thus: "Sure, there's an epistemic norm according to which I should undertake the inquiry that you're suggesting, since I have good reason to believe that doing so would help me get to the truth regarding the question of God's existence. But look! I have overriding reasons not to undertake that inquiry. I want to be able to continue telling people I meet at parties that

I'm an agnostic, since people think that agnostics are deep, thoughtful, and intellectually moderate."

You ought not, I submit, let her get away with that. You ought not let her get away with it even if she's being quite sincere when she says she prefers maintaining that image over getting to the truth on the question of God's existence.

So the overriding reasons we're talking about need some objective meat on their bones. The epistemic agent's mere subjective preference for some other good over getting to the truth and avoiding error is not sufficient. There needs to be some objective sense in which the thing that the epistemic agent is preferring over the epistemic goal actually should be so preferred.

There is an important caveat to that objectivity requirement that should be unsurprising given what's been said previously. We can conceive of agents who blamelessly but mistakenly believe that they have good overriding reasons to violate some epistemic norm. Imagine Tim, who has been raised his entire life in a hyper-fundamentalist religious community that insists that to inquire into the truth or falsity of its doctrines is to incur divine wrath. Suppose that all the evidence he has – the testimony of his parents and religious leaders, religious experiences he's had that seem to corroborate the authority of his religious leaders, etc – supports the belief that his community is reliable when it says this. Tim might quite blamelessly believe that he that he has overriding reason to, in this case, violate the epistemic norm that one ought inquire into the truth or falsity of one's beliefs. He would be mistaken, but to blame him for violating the norm would be to ignore the importance of the *mens rea* requirement for

blame. He has no way of knowing that what he thinks are overriding reasons to violate the norm are not actually so.

Return to Sally the agnostic. (And assume that she does not have some strange set of evidence that leaves her blameless in believing that appearing interesting at parties is more important than getting to the truth on the matter of whether God exists.) We blame her, ultimately, because she is violating an epistemic norm for frivolous reasons. The relative weights that she assigns to the epistemic goal and her goal of seeming interesting at parties are simply absurd. The epistemic goal is not all-important such that it always overrides every other goal, but surely it is very important. Sally, however, is treating it as though it were a trifle. She has adopted a fundamentally unserious stance towards it.

That reveals something important about epistemic blame. Epistemic blame involves a judgment about the agent's attitude towards the epistemic goal. If an agent is aware of an epistemic norm (which is required for the *mens rea* component) but violates it anyway with no overriding reason, they are not giving the goal that governs that norm its due. They are, to some degree, setting aside the epistemic goal and conducting at least part of their epistemic lives in a way that is perversely disconnected from that goal.

CHAPTER THREE

Epistemic and Moral Norms

Introduction

So far, we have been proceeding as though epistemic blame is a fundamental kind of blame – that it is not a subspecies of some other species of blame. On our account of blame, that amounts to proceeding as though epistemic norms are not a subspecies of some other species of norm. There is, however, a long tradition in Western thought of moralizing the cognitive life by treating epistemic concepts as moral concepts, and this tradition has found contemporary defenders.¹

The question of whether epistemic blame is a species of moral blame needs to be distinguished from a variety of related questions about the relationship between epistemic and moral concepts. Our question is not whether epistemic evaluation is, in general, moral evaluation. That question can be straightforwardly answered in the negative. There are all kinds of evaluations that deserve to be called epistemic – evaluations of someone's intelligence, of the reliability of their cognitive equipment, etc. – that are obviously not moral evaluations.

Nor is our question whether epistemic justification, understood as a normative component of knowledge, is a kind of moral justification. It clearly is not on many popular views of the normative component of knowledge – externalist ones, for instance. More generally, on any view on which it is possible for someone to lack the normative

¹ See, for instance, Linda Trinkaus Zagzebski, *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*, (Cambridge: Cambridge University Press, 1996.)

component of knowledge through no fault of their own, that component has nothing to do with moral justification.

Our specific question has received less attention in the literature than have other questions about the relationship between epistemic and moral concepts. Discussion of related questions must be adapted to ours. If a duty is something one can be blamed for not fulfilling, then the question of whether epistemic duties are moral duties is particularly well-suited to such adaptation.

In this chapter I will examine and reject arguments from W. K. Clifford and John Locke for treating epistemic blame as a species of moral blame. I will then criticize some arguments against doing so. I will close by giving some reasons for thinking that epistemic norms are in fact moral norms, and that epistemic blame is therefore a kind of moral blame.

Clifford and Locke

Clifford

Clifford, in his classic essay “The Ethics of Belief,” holds that we have duties related to our cognitive lives. These duties are clearly moral ones. Those who fail in their duty are said to be “no longer to be counted as honourable men,”² and to have incurred “a stain which can never be wiped away.”³ These duties are to believe only on sufficient evidence and to engage in thorough and responsible inquiry – in other words, to fulfill the evidentialist and responsibilist norms.

² W. K. Clifford, *The Ethics of Belief and Other Essays*, (Amherst, NY: Prometheus, 1999) 72.

³ *Ibid* 74.

To whom do we have these duties? To other people. Why do we have these duties? For two reasons. First, our beliefs lead us to act or to fail to act in various ways. So, if we act on beliefs that we have not arrived at through rigorous investigation, we risk acting on beliefs that are false. To risk that is to risk acting in ways that could harm others – Clifford’s famous case of the ship owner who sends an unseaworthy ship to sea filled with passengers because he has convinced himself that it is seaworthy despite the evidence to the contrary illustrates this. But to risk harming others in this way is immoral. It remains immoral even if the belief that has been unrigorously arrived at happens, by chance, to be true, for exactly the same reason that my driving while intoxicated would be immoral, even if I happened by chance not to cause an accident.

The second reason I have a duty to believe only after rigorous investigation is that if I do not, I will contribute to a broader culture of epistemic misbehavior. Even if there’s no reasonable way a belief of mine could lead to actions that might harm another – suppose I believe in invisible lions who demand that I behave as though they did not exist, for instance – in believing without rigorous investigation I become credulous. If I am credulous, then I am all the more likely to believe false things in the future, with all the negative results that might come of that. And what’s more, I help make others credulous as well. I pass on my low epistemic standards to my children. I reaffirm my neighbor in his bigotry and superstition. I help create a culture that is simply not concerned with the truth, with all the disastrous consequences that follow from that.

One objection to this line of thought is that even if every instance in which someone fails to believe on the basis of evidence after rigorous inquiry is morally objectionable for precisely the reasons that Clifford describes, that does not demonstrate

the epistemic norms are moral norms, nor that epistemic blame is moral blame. After all, the same action can violate more than one norm. If I murder my opponent in a racquetball game to prevent him from returning my serve, I have violated moral norms, legal norms, and the norms of racquetball.

Why not say that something similar is happening when people violate epistemic norms? The act of believing that violates the epistemic norms will always also violate moral norms, because they will always involve a sort of carelessness with the fates of others. But the carelessness with the fates of others is a separate problem. The violations of the epistemic norms, considered in themselves, are just that and nothing more. It is not that the epistemic norms are moral norms. It is simply that any act that violates epistemic norms will also violate entirely separate moral norms.

One reason not to go this route is that we seem quite capable of recognizing when one act violates more than one norm, and of deploying the language of blame in a way that calls attention to that fact. When someone commits a crime, we often say things like “That was illegal, and it was also wrong.” We quite explicitly blame the criminal twice, under two different kinds of norms. There does not seem to be anything parallel going on with the way we talk about violations of epistemic norms.

A more fundamental problem with using reasoning of the sort the Clifford gives us to support the claim that epistemic norms are a kind of moral norm deals with the explicitly consequentialist nature of Clifford’s moral argument. We should be careful to abide by the epistemic norms because of the bad consequences, in the short or long term, that will follow if we don’t.

But sometimes I can quite scrupulously follow my epistemic duties while pursuing projects that are harmful to others. My local shaman tells me that this spell will help me kill those who have insulted me. But should I just take his word for it? Sure, he's the shaman, but that does not free me from my obligation to investigate the efficacy of magic spells for myself. After rigorous investigation I conclude that magic spells are not an effective means of having my vengeance, and decide to just poison my enemies instead.

Here I've done my epistemic duty quite nicely. Blindly taking the shaman's advice would have violated some epistemic norms. As far as consequentialist reasoning goes, though, it certainly seems better for me to trust the shaman and fail to kill my enemies than to critically evaluate his plan, come up with a better one, and succeed at killing my enemies. This remains true even if we include the harm done by the reinforcing of my tendency towards credulity and the risk of that credulity coming to be reflected in my society. Surely one person lazily accepting the word of a shaman one time won't do enough damage on that score to equal the harm done by the murders of several people. So it is difficult to see how we could give a Clifford-style consequentialist argument for the immorality of my blindly trusting the shaman. But it would still violate an epistemic norm.

Or suppose that God appears to me and presents me an offer. "Believe that the number of stars is even," he says. "If you do, you will be immediately taken up into Heaven, where you will enjoy eternal bliss and your fellow humans will be entirely insulated from the consequences of the bad epistemic character that you will have acquired. None of them will know what you've done, so they won't be tempted to follow

your example. Furthermore, I'll end all war, disease, and famine, and fill all humankind with a passionate love of truth and with skill and tenacity in obtaining it."

Now, I have no evidence that the number of stars is even.⁴ So I can't believe it without violating an epistemic norm. But again, the moral arguments Clifford gives against believing on inadequate evidence seem entirely toothless here. Indeed, I seem to have excellent consequentialist reasons for accepting God's offer. That does not change the fact that I am violating an epistemic norm. I may not be *blameworthy*, since my excellent non-epistemic reasons for violating the norm may mean that I'm not culpable in doing so, but the evidentialist norm has clearly been violated.

So there are possible violations of epistemic norms that, at least as far as Clifford-style consequentialist arguments can tell, are not violations of moral norms. So we should not appeal to such arguments to show that epistemic norms are moral norms. We therefore should not appeal to them to show that epistemic blame is a kind of moral blame.

Locke

John Locke, in his *Essay Concerning Human Understanding*, offers a theological argument for the immorality of misusing our rational faculties. Our moral duty to abide by the epistemic norms is, for Locke, a duty owed to God:

He that believes without having any reason for believing, may be in love with his own fancies, but neither seeks truth as he ought, nor pays the obedience due to his Maker, who would have him use those discerning faculties he has given him, to keep him out of mistake and error. [...] He that does otherwise [than use his rational faculties properly] transgresses against his own light, and misuses those

⁴ Suppose that I am, for whatever reason, incapable of inferring P from the fact that God has encouraged me to believe P.

faculties which were given him to no other end, but to search and follow the clearer evidence and greater probability.⁵

So, I have a moral duty to use my rational faculties properly because God gave them to me and wants me to use them in that way, to achieve the end for which he gave them to me.

There are at least two different ways to understand this line of argument. We could understand it as claiming that God made our cognitive faculties with certain ends in mind, that those ends determine the purpose of those faculties in the same way that the purpose of any artifact is determined by the intent of its creator, and that the purpose of our cognitive faculties places moral constraints upon their use. So we should use our cognitive faculties in a way that will help us get to the truth because that's the end for which God made them, which means that's what they're for, and we are morally obligated to use our cognitive faculties in a manner consistent with what they are for.

There is language in the passage quoted above to support such a reading. The phrase "misuses those faculties which were given him to no other end" could certainly be read in a way that is friendly to such an interpretation. But it needn't be. After all, the passage claims that our cognitive faculties were "given" for a particular purpose, not that they were made for a particular purpose. That is an important distinction. If I receive something as a gift, I may have an obligation to use it in a manner that is consistent with the purposes of the one who gave it to me, even if I have no general obligation to use things in a manner consistent with their purpose as determined by the intentions of their makers.

⁵ John Locke, *An Essay Concerning Human Understanding* ed. Peter H Nidditch. (Oxford: Clarendon Press, 1975), 687-688

That brings us to the other possible interpretation of Locke's line of argument – I have a moral obligation to use my cognitive faculties properly not because of their purpose as set by God's intentions, but because God has told me to and God's instructions are, at least in this case, sufficient to generate moral obligations. Locke's claim that the epistemic malefactor fails to render "the obedience due to his Maker" certainly supports that reading, and it fits well with Locke's broader understanding of moral obligation.

There are a number of objections to this line of argument. One is that it presupposes a divine command theory of moral obligation, and divine command theories of moral obligation are generally implausible.

It is true that Locke's argument here fits within the context of a broader divine command theory. But it is false that the claim that we are morally obligated to use our rational faculties properly because God told us to presupposes a divine command theory of moral obligation. That is because the claim that commands can generate moral obligations is consistent with a wide range of views of the nature of moral obligation. Divine commands can be a *source* of moral obligations without it being the case that to be a moral obligation just is to be a divine command. To take one example, a utilitarian about moral obligation could easily hold that I am morally obligated to do what God tells me to do, simply because a general policy of obeying divine commands does better in the utilitarian calculus than the alternatives.

We are not committed to a parental command theory or a governmental command theory of moral obligation when we hold that parental and governmental commands can morally oblige. That is because, regardless of our view of moral obligation, we think that

people who occupy certain roles relative to us have binding authority over us, at least in certain spheres. So, what role might we say that God plays, relative to us, by virtue of which his commands morally oblige us within the sphere of the use of our cognitive faculties?

Well, we could say that he fulfills the role of God. Perhaps God, simply by virtue of being God, has the authority to issue binding commands regarding the use of my cognitive faculties. For God to have such wide-ranging authority simply because he is God does not, again, commit us to a divine command theory since, as mentioned above, a utilitarian could easily hold that God has such authority.

Perhaps he has the authority to command because he fulfills the role our creator. Because God is responsible for our existence, he has the authority to give us binding commands. Perhaps this flows from a more general obligation of gratitude, or from some general obligation similar to the one that obliges us to obey our parents. Locke's use of "Maker" language supports such a reading.

Maybe he has the authority to command us to use our cognitive faculties in particular ways because he is the one who gave them to us. It is, after all, a plausible enough general principle that the giver of a gift has the authority to place binding constraints on the use of that gift. Locke's claim, which is repeated twice in the passage quoted above, that the epistemic malefactor has misused faculties that were "given him" might suggest that this is what Locke has in mind.

There are, then, plenty of ways to understand the authority that Locke is ascribing to God here that do not presuppose a divine command theory of moral obligation. So worries that we might have about the plausibility of divine command theories should not

give us pause when it comes to the claim that we have a moral obligation to fulfill the epistemic norms that is grounded in God's commands.

Another worry about Locke's approach is identical to one that came up in the discussion of Clifford. Suppose that Locke is right that we have a moral obligation to fulfill the epistemic norms, and that that obligation is grounded in God's commands. How does that show that epistemic norms are moral norms? God could command me to obey the law, and that command could generate a moral obligation to obey the law. The same criminal act, then, could violate both a moral norm and a legal norm. But that wouldn't mean that legal norms are moral norms. The legal norms are what they are because of facts about what the law is, not because God has commanded me to obey the law. So, when I fail to fulfill an epistemic norm, why not just say that the same act has also violated an entirely separate moral norm, one grounded in the commands of God?

This is a genuine worry about using Locke's line of argument to show that epistemic norms are a kind of moral norm. But the same response that was given in the discussion of Clifford has the same force here that it did there – we are pretty good at detecting when multiple norms are being violated, and we do not seem to detect the violation of multiple norms when considering acts that violate epistemic norms, at least not generally.⁶

The final worry about this line of argument is parallel to another worry about Clifford's. Locke's account is based on commands, and commands are contingent things. They can be undone. So imagine God said to me "Despite the fact that you have no

⁶ Sometimes we might. We might say that the ship owner in the Clifford case violated both an epistemic and a separate moral norm. But that is consistent with saying that epistemic norms are a kind of moral norm. When we make that judgment about Clifford's ship owner, we might simply mean that he violated two different moral norms – the moral norms governing epistemic behavior and the moral norms that say we should exercise reasonable caution to avoid acting in ways that will harm other people.

evidence regarding the number of stars, you have my permission to believe either that it is even or that it is odd, as may happen to suit your fancy. Any general command I may have given that says you should only believe on sufficient evidence is suspended in this case.”

If that happened, then there would be no Locke-style argument to the conclusion that I have a moral duty not to believe that the number of stars is odd unless I have good evidence that it is. Nevertheless, it would still violate an epistemic norm. God’s giving me permission to believe without sufficient evidence does not change the fact it is an extremely poor way to go about getting to the truth and avoiding error, which is the epistemic goal that governs the epistemic norms. It is still bad epistemic practice, even if God indicates that my bad epistemic practice is, at least in this instance, fine with him.

So, if I went ahead and believed that the number of stars was odd in this case, I would be violating an epistemic norm but not, as far as Locke-style arguments can tell us, a moral norm. Locke style arguments, then, are not a promising way of showing that epistemic norms are moral norms.

Arguments against Identifying Epistemic and Moral Norms

The modern literature contains two broad kinds of arguments against the claim that epistemic norms are moral norms. The first has to do alleged cases in which moral and epistemic obligations come apart – such that the epistemically right thing to do is the morally wrong thing to do, or the epistemically wrong thing to do is the morally right thing to do. The second holds that what we do epistemically is not within our control in the right sort of way to be the proper subject of moral evaluation.

The Argument from Doxastic Involuntarism

Here, in a rough, initial way, is the argument from doxastic involuntarism against identifying epistemic blame as a kind of moral blame. We do not hold people morally blameworthy for things that we acknowledge are beyond their control. That is because we think that ought, for any sense of “ought” on which failing to do as we ought is something for which we are properly subject to moral blame, implies can. So if we are going to blame people morally for violating epistemic norms, then it has to be the case that they could have avoided violating them.

But one of the epistemic norms is the evidentialist norm. The epistemic norm tells us that we should hold beliefs that fit our evidence, and refrain from holding beliefs that do not fit our evidence. That means that blaming someone morally for violating the evidentialist norm presupposes *doxastic voluntarism* – the thesis that holding or refraining from holding a particular belief is subject to our direct voluntary control. But doxastic voluntarism is false. We do not decide to believe or refrain from believing. We simply find ourselves either believing or not believing. That means that we should not blame people morally for violating the evidentialist norm – blame for violating that norm cannot be moral blame.

There are three things that can be said in response to the argument from doxastic involuntarism. First, the objection only touches the evidentialist norm. Perhaps I do not have the right kind of control over my believings to make it appropriate to morally blame me for a particular belief that fails to match my evidence. But I certainly have all kinds of control over my processes of inquiry – reading a newspaper, performing an experiment, or intentionally forgoing either are as unproblematically voluntary as giving to charity or

investing in a company that I know to have atrocious labor practices. So if any standard human action is sufficiently voluntary to be a subject of moral evaluation, then so are the sorts of actions that fall under the responsibilist norm.

Now, if it is true that the doxastic involuntarism argument shows that the evidentialist norms are not moral norms, then it is in general false that epistemic norms are moral norms. But it could still be the case that *some* epistemic norms are moral norms. And if, as I intend to argue later, our actual epistemic evaluations of others should focus entirely on the responsibilist norms, then it could still be the case that all proper epistemic evaluation is moral evaluation.

Second, when we judge that someone was not sufficiently free in performing an action that violates a norm, that judgment affects our assignment of culpability, not our judgment that a norm has been violated. If I can't help but believe something that does not fit my evidence, then it is not my fault, and I thus should not be blamed for it. But that does not mean that there is no norm that says I should believe only that which fits my evidence.

What the doxastic involuntarism argument shows, if anything, is simply that we are across the board not culpable for violating the evidentialist norm, and thus that we should not be blamed for it. It does not show that the evidentialist norm is not a moral norm.

The third point is that doxastic involuntarism, if it is a threat to understanding epistemic blame as a kind of moral blame, is not *just* a threat to understanding epistemic blame as a kind of moral blame. The right kind of voluntariness is necessary for culpability, and culpability is a necessary feature of all blame, not just moral blame. So if

I do not have the kind of control over whether or not I violate the evidentialist norm that I would need in order to be blameworthy for doing so, then I shouldn't be blamed epistemically any more than I should be blamed morally. So if doxastic involuntarism is a threat to epistemic blame considered as moral blame, it is also a threat to epistemic blame *simpliciter*.

In general, the culpability requirement for blame means that the argument from doxastic involuntarism is poorly suited to a discussion of the relationship between epistemic and moral blame. To see why, consider this version of the argument from doxastic involuntarism by Susan Haack, directed against the claim that epistemic justification is a moral concept:

A person is epistemically unjustified in believing that *p* just in case his evidence isn't good enough. But he can't be morally at fault in believing that *p* unless his belief is willfully induced. And his evidence may not be good enough even in cases where his belief is not willfully induced. So it is possible that there should be cases where a person is epistemically unjustified but not morally at fault.⁷

Even if doxastic involuntarism is false as a general claim about beliefs, it is still clearly conceivable that one might have some beliefs that are not willfully chosen. I could be brainwashed, or have a belief directly implanted in my mind by God. Such beliefs may not fit my evidence. They would thus be epistemically unjustified, but I certainly would not be subject to moral reproach for them.

But these unjustified but morally blameless beliefs are morally blameless because they are not culpable. The involuntariness of the beliefs is incompatible with their culpability. But if such beliefs are inculpable, then they are not epistemically blameworthy, since culpability is a necessary condition for blameworthiness. Any belief

⁷ Susan Haack, "The Ethics of Belief Reconsidered," in *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue* ed. Matthias Steup (Oxford: OUP, 2001) 21-30

that is morally blameless because it is involuntary is also epistemically blameless. We thus cannot appeal to the possibility of involuntary beliefs that do not fit our evidence to show that there can be beliefs that are epistemically but not morally blameworthy, because no such belief is epistemically blameworthy either.

Worries arising from doxastic involuntarism, then, do not give us a good reason to deny that epistemic blame is moral blame.

Do Moral Epistemic and Moral Blame Come Apart?

If epistemic blame is a species of moral blame, then we should expect there to be no cases in which the two come apart – in which we would blame an agent epistemically but praise her morally, or *vice versa*. But there cases to be found in the literature in which, at least at first glance, it appears that that is precisely what we should do.

The trusting husband. Take, for instance, the case of a husband who has evidence of his wife's adultery, but does not believe that she is having an affair because of the trust that is an appropriate part of his relationship with her.⁸ He is violating the evidentialist norm by not believing according to his evidence, but in at least some cases we would not be inclined to morally blame him. We would be inclined to say that letting his trust in his wife rather than his evidence guide his beliefs regarding her fidelity is praiseworthy, if not obligatory. One is supposed to trust one's spouse, after all. But if the man is morally praiseworthy despite being epistemically blameworthy, then epistemic blame is not a species of moral blame.

⁸ This case is mentioned, among other places, in Richard Fumerton, "Epistemic Justification and Normativity," in *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue* ed. Matthias Steup (Oxford: OUP, 2001) 49-58

There are several ways to respond to this argument. One obvious way of doing that is to point out that we can quite coherently say that someone is all-things-considered morally praiseworthy in some situation while still holding that they are morally blameworthy in some respect. Consider someone who accepts an absolute prohibition against lying. Such a person might respond to the standard “lying to the Nazis about whether you’re hiding Jews” objection to this position by saying that a person who lies in such a case is (a) morally praiseworthy for standing against injustice and protecting the life of another at great personal risk, and (b) morally blameworthy for lying. They may further hold that the moral praise for (a) greatly outweighs the moral blame for (b), so that the person is, all things considered, morally praiseworthy with respect to the Nazis-at-the-door situation.

Likewise, we can make sense of our intuition that the husband in the case described above is morally praiseworthy even on the assumption that epistemic blame is a species of moral blame. Trusting one’s spouse is morally praiseworthy. Failing to believe in accordance with one’s evidence is morally blameworthy. When one trusts one’s spouse by failing to believe in accordance with one’s evidence, one has done something that is morally praiseworthy in one way and morally blameworthy in another. If one should be praised for trusting more than they should be blamed for not believing in accordance with the evidence – which is plausible enough, given that marital trust involves an interpersonal relationship of first importance – then one is, all things considered, praiseworthy for believing as one does.

There is at least one problem with this way of responding to the case. The husband in this case is not being praised simply for believing that his wife was not having

an affair. Countless people do that every day without accruing any particular merit. He is being praised for believing that his wife is not having an affair *despite the fact that he has good evidence that she is*. It is the presence of evidence to the contrary that makes his belief meritorious. So is not, as in the Nazi case, that he acted for the sake of a praiseworthy goal in a blameworthy manner. The very same act under the very same description – believing that his wife is faithful despite the fact that he has good evidence that she is not is, if epistemic blame is a species of moral blame, both morally praiseworthy and morally blameworthy. That is a consequence that we should want to avoid. So another way of responding to the case needs to be found.

Another way of dealing with the case appeals, yet again, to the culpability requirement in our account of epistemic blame. If the man believes that his wife is faithful despite the fact that his best evidence suggests that she is not, he has violated an epistemic norm. But to be blameworthy, he must have violated the norm *culpably*. If he has overriding reasons to believe that his wife is faithful – moral reasons, relational reasons, etc. then we would not treat him as culpable for violating the evidentialist norm. If the moral praiseworthiness of trusting his wife gives him an overriding reason to do so, then, we he is not epistemically blameworthy in doing so. So this is not a case in which we blame him epistemically but don't blame him morally, because it is not a case in which we blame him epistemically.

This reply is problematic because, while it defuses the case as a counterexample to the claim that epistemic blame is moral blame, it does nothing to prevent the case from showing that epistemic norms are not moral norms. Even if the man in the case is not culpable, he still seems to have violated an epistemic norm. If our intuitions in the case

are that he has not violated any moral norms, then that would suggest that epistemic norms are not moral norms.

Given our account of blame, the kind of blame a given instance of blame is depends upon the kind of norm that is judged to have been violated. So if epistemic norms aren't moral norms, then epistemic blame is not moral blame.

The best way of responding to this case involves pushing back against the description of the case. We are asked to consider the man's evidence as something opposed to his trust in his wife, and that all of his evidence is therefore on the side of his wife's being adulterous – his trust in her has nothing to do with evidence. But his trust in her is a belief about her character, or a disposition to form certain beliefs about her character, and as such it is responsive to evidence. There is thus no obvious reason not to count his trust in her as part of the evidence that he has to bring to bear on the question of her fidelity.

So, either the husband's trust in his wife is reasonable, given his evidence, or it is not. Suppose that it is. Suppose that the intimate knowledge of her character that he has gained over the course of their marriage is such that it gives him excellent grounds for trusting in her fidelity. Then that knowledge is evidence, and he can and should account for it when forming beliefs about whether or not his wife is having an affair.

Treating the man's trust in his wife in evidential terms does not reduce the moral good we see in his trust to the moral good of believing on the basis of evidence.

Attending properly to certain kinds of evidence in certain situations can be difficult, and requires strength of moral character. A less virtuous person in his situation might have been inclined to ignore the evidence he had about his wife's character, and to focus

exclusively on the evidence that seems to indicate that she is unfaithful. A less virtuous person might be given to jealousy, or enjoy believing the worst about those closest to him, or believe that being able to play the part of the betrayed spouse will give him more power in the relationship. Saying that, in trusting his wife, the man is simply following his evidence does not threaten the non-epistemic moral approval we seem inclined to give him.

Or, suppose that the man's trust in his wife is not reasonable, given his evidence. Suppose that he has no intimate knowledge of her character that gives him reason to trust her not to cheat on him, either because he has no intimate knowledge of her character or because the intimate knowledge of her character that he does have does not give him reason to trust her. Suppose he knows for a fact that she has cheated on him before. Suppose she has explicitly told him that she is not sorry for doing so, and would do so again should a suitable opportunity arise. So, in addition to whatever external evidence he has of her unfaithfulness – suspicious phone calls, implausible business trips, and so forth – the knowledge he has of her character also tends to support the conclusion that she is having an affair.

I think that any moral approval that we would be inclined to give to the man's believing his wife to be faithful vanishes when the case is filled out in this way. He should not trust her – whatever sort of trust might be proper to marriage, it is not trust that is as willfully blind to the evidence as this.

So, depending on how we fill out the details of the case, either the man is not in fact epistemically blameworthy in believing his wife to be faithful, or we should not give

a positive moral evaluation to his doing so. Either way, the case is not one in which epistemic blame and moral blame come apart.

The optimistic patient. Another case that might be used to argue against identifying epistemic blame as a species of moral blame is one given by Roderick Firth:

Now let us suppose that John Doe is suffering from an illness that is usually fatal, but believes with deep conviction that he will recover. The fact that John has this optimistic belief might actually contribute to his recovery. Or at least it might make him more cheerful during his dying days, which in turn might ease the pain of others who are close to him. In either case it would be a virtue or *merit* of John's belief that it has the good consequences that it does for himself or others. [...] Merit of this kind might also be classified as ethical.⁹

The merit might be classified as ethical if it results from an effort on John's part. There are kinds of ways that John could, at least indirectly, try to cause himself to believe that he will recover. He might constantly remind himself of research suggesting that a cure for his condition is right around the corner, while ignoring subsequent research calling that finding into question. He might take the fact that all his doctors have confident mannerisms and wear well-ironed shirts as evidence that they are extremely competent, while ignoring other evidence that they are, in fact, rather incompetent.

To help fit the case to our question, suppose that John, due to his importance to his family or to humanity at large, has a duty to live. In believing that he will recover he is fulfilling that duty. If he did not believe he would recover (and was aware that he could get himself to believe that he'd recover if he tried, and knew that believing that would help him recover) he would be neglecting his duty. In other words, he would be morally blameworthy.

⁹ Roderick Firth, "Epistemic merit, intrinsic and instrumental." *The American Philosophical Association Centennial Series* (2013): 5-18.

But in causing himself to believe this while knowing full well that it did not match his evidence, he would be freely and deliberately violating epistemic norms – he would be epistemically blameworthy. So he would be epistemically but not morally blameworthy. Epistemic and moral blameworthiness would not go together in this case.

One response would be to point out that the duties we have depends on the means we have to accomplish certain goals. I have a duty to provide for my children. But that just means that I have a duty to employ a certain range of means to provide for my children. I could provide for my children by wiping out the entire population of Ohio and claiming all of its resources for my own, but I don't have a duty to do that. I have a duty not to do that.

It is very plausible that nobody can ever have a moral duty to employ a morally wrong means to achieve any end whatsoever. If so, John cannot have a moral duty to believe that he will live as a means to achieving the end of his survival if believing he will live is, under the circumstances, morally wrong. But if epistemic blame is moral blame then, under the circumstances described, believing that he will live is morally blameworthy, and thus morally wrong. That means that anyone who believes that epistemic blame is moral blame should not accept the possibility of the case as described to begin with.

Why Think that Epistemic Blame is Moral Blame?

So, neither the argument from doxastic involuntarism nor the alleged cases in which moral and epistemic blame come apart give us a good reason to deny that epistemic blame is a species of moral blame. But the arguments derived from Clifford and Locke at the beginning of the chapter for believing that epistemic blame *is* a species

of moral blame are fairly weak. So what can be said in favor of the thesis that I have been defending?

The Inescapability of Moral and Epistemic Evaluation

We cannot escape moral evaluation. If I violate some moral norm and you blame me for it, I cannot respond to your blame by telling you that I have no interest in obeying moral norms. If anything, that would only make you blame me more. Moral norms are absolute in that I cannot exempt myself from being evaluated under them by disassociating myself from the goods at which those norms aim. I *ought* to be concerned with those norms, no matter what. If I am not, that does not excuse me from blame, but rather increases it.

Not all norms are like that. The norms related to various activities that we might call practical norms are not. If I intend to be a curler, then there are all kinds of things I ought to do. I ought to learn all I can about curling strategy. I ought to practice my delivery and my sweeping. These are norms of curling. If I don't fulfill them, I am properly subject to blame. But I am properly subject to blame only if I have associated myself with the goal that governs those norms – the goal of curling. It would be rather pathological of you to blame me for not practicing my delivery and sweeping if I had no interest whatsoever in curling.

Epistemic norms seem to be like moral norms and unlike curling norms in this respect. I cannot exempt myself from blame under the epistemic norms by disassociating myself from the goal of getting to the truth and avoiding error. If I really have no interest in that goal, that just makes things worse, not better.

That moral norms have this absolute quality is an essential part of their nature. That's just what morality is about – what we ought to do no matter what, regardless of our interests, desires, etc. That epistemic norms are also about what we ought to do no matter what, regardless of our interests, desires, etc. is indicative.

Simplicity

It becomes more indicative when we bring in considerations of simplicity. There is some advantage to simplicity in philosophical theories. All things being equal, it is an advantage of a theory of the relationship between phenomena of type A and phenomena of type B if it accounts for one type in terms of the other rather than leaving both as fundamental, irreducible types of phenomena. So, all things being equal, it would be better if we could say that either epistemic or moral blame is a species of the other, which would mean that either epistemic or moral norms are a species of the other.

But, despite the fact that it is related to some ideas that have respectable and ancient roots, the idea that moral blame is a species of epistemic blame is a non-starter. We can act in a morally blameworthy way even if we are epistemically blameless in all relevant respects. We should not analyze the moral blameworthiness of a bad act in terms of epistemic blameworthiness that leads to some mistaken belief about the moral qualities of that act, because no such epistemic blameworthiness need exist. *Akrasia* is a real phenomenon.

So if we can explain either of these two kinds of blame in terms of the other, we will be explaining epistemic blame in terms of moral blame. And simplicity suggests that we should want to. We have already seen that the arguments against doing so are unsuccessful. And the similarities between the two noted above make the simplicity

considerations starker. Both, after all, deal with blame under inescapable norms governing human behavior. Why would we want to posit a fundamental category of normativity for the inescapable norms governing human cognitive activity, any more than we would for the inescapable norms governing any other sphere of human activity? There is no theoretical gain to doing so. So we should not do so.

This does not mean that there is no category of epistemic normativity that cannot be reduced to moral normativity or any other kind of normativity. Perhaps justification, as knowledge theorists are interested in the term, constitutes such a category.

Escapable Epistemic Blame

The above proceeds on the assumption that in all cases I ought to be concerned with getting to the truth and avoiding error, so that I cannot avoid blame simply by disassociating myself from that goal. But with respect to some truths, that is plausibly not the case. Perhaps there are some truths that human being can quite legitimately not care about acquiring. How many grains of sand there are in this heap can serve as a plausible example. I might still chose to form a belief on this question. And I might do so in accordance with my evidence or not, after appropriate investigation or not. If I fail to believe according to my evidence, or form a belief without investigation properly, you might blame me. But it is not moral blame, since I can avoid it simply by not caring about the goal of getting to the truth regarding the question.

But it is not thereby an irreducably epistemic kind of blame. All you are doing is blaming me for failing to take the appropriate actions to attain a goal -- a true belief about

the number of grains of sand in the heap -- that I claim to want to attain. My failure then is simply one of practical rationality.¹⁰

So, to sum up, we have seen that the arguments against identifying epistemic blame with moral blame fail. While arguments drawn from Locke and Clifford in favor of doing so likewise fail, considerations of simplicity give us reason to do so, at least with respect to truths that we ought to care about. And violations of epistemic norms with respect to truths we optionally care about are simply violations of practical norms. In neither case, then, is epistemic blame a fundamental kind of blame.

I take it that we do not much care about epistemic blame regarding optional truths like the number of grains of sand in a heap. So I also take it that all the instances of epistemic blame that we care about are instances of moral blame.

¹⁰ For an argument for reducing failures under the responsibilist norm to either moral failures or failures of practical rationality, see Trent Dougherty, "Reducing Responsibility: An Evidentialist Account of Epistemic Blame." *European Journal of Philosophy* vol 20 no 4 (2012): 534-547.

CHAPTER FOUR

Blaming Well is Hard

Introduction

If blame is a kind of judgment, then questions of epistemic normativity arise in relation to it. Instances of blame could be good or bad in all the ways other judgments could be good or bad. They could be true or false, well-grounded or not well-grounded, reliably formed or unreliably formed, blameworthy or blameless.

As with all of the other judgments that we make, then, we should be concerned with making the judgments involved in epistemic blame well.¹ That means paying careful attention to the limits of our cognitive access to the factors involved in such judgments. In this chapter, I will argue that our cognitive access to the factors relevant to epistemic blame is very limited. We are not very good at knowing whether someone has culpably violated either the evidentialist or the responsibilist norm. And many cases that are treated both by philosophers and the public at large as clear-cut instances of epistemic blameworthiness are nothing of the sort.

The Evidentialist Norm

To determine whether someone is epistemically blameworthy because they have culpably violated the epistemic norm, we must first determine whether they have violated an epistemic norm at all. We must determine whether or not the belief in question fits the evidence. At least in normal cases, that will involve determining what the evidence is.

¹ We have at least the same reasons to make such judgments well as we do for any other judgment, but there are quite plausibly additional moral reasons as well.

Evidence as Publically Available Data

There are views of evidence on which that is a fairly straightforward task, since the evidence is public and unproblematically open to our examination. Take, for instance, this discussion by Richard Yetter-Chappell on the evidential credentials of theism:

Cosmological and fine-tuning arguments don't support anything stronger than minimal deism. The ontological argument is obvious sophistry, and the modal version is question-begging. The design argument was defeated by Darwin; in its place we have the problems of evil and divine silence: the world just doesn't look anything like we'd expect it to if overseen by an omni-max god. In short, it seems to me, we seem to have every reason to reject theism, and no good reason to accept it.²

Yetter-Chappell is not directly addressing the *blameworthiness* of religious belief, but rather its unreasonableness. Nevertheless, he is clearly claiming that religious beliefs (more specifically, theism,) do not fit the evidence.

What is the evidence? The standard repertoire of theistic and atheistic arguments. Those arguments are publicly available. We can examine the merits of those arguments, and in doing so come to know what attitude towards theism they support. If someone takes an attitude towards theism other than the one supported by that evidence, then their attitude does not fit the evidence. As far as determining blameworthiness goes, all that is left is determining culpability.

A similar view is endorsed by the case given in the preceding chapter of the man who trusts his wife despite apparent evidence that she is having an affair. The case supposes that he is epistemically unjustified in believing his wife to be faithful, because that belief does not match the evidence that he has. That evidence consists, again, of publicly available facts about his wife's behavior. It is assumed that the man is in the

² Richard Yetter-Chappell. "Is Religious Belief Reasonable?" <http://www.philosophyetc.net/2013/10/is-religious-belief-reasonable.html> October 19, 2013

same epistemic position with respect to his wife's fidelity as any third party aware of the same facts about her behavior would be. Any private, first-person factors that the man might bring to bear on the question do not count as evidence.

A similar dynamic is quite common in non-academic discourse, particularly involving contentious moral, religious, or political issues. Some set of publicly-available evidence is surveyed and declared to be "the evidence," and the rationality of those whose beliefs do not appear to fit that evidence is impugned.

There are very good reasons not to adopt a policy of blaming people epistemically when they fail to believe in a way that matches the evidence if "evidence" is to be understood in this way. People could quite easily fail to believe in accordance with the publicly available data without being in any way blameworthy. For starters, people could be unaware of some of the publicly available data.³ The theist may simply have never heard a compelling presentation of the argument from evil or from divine silence. The trusting husband may simply fail to notice his wife's suspicious behavior. In neither case would we want to blame them under the evidentialist norm for not believing in accordance with the publicly available data.⁴

There are two ways to avoid holding people blameworthy under the evidentialist norm when they fail to believe in accordance with publically available data of which they are unaware. First, we could say that epistemic agents who do are unaware of relevant publically available data are exempt from blame not because their beliefs in fact fit the evidence, but because, since they were ignorant of some of the evidence, they are not

³ Yetter-Chappell acknowledges this point by limiting his discussion of the rationality of theism to "well-informed agents."

⁴ Although we may, under certain circumstances, blame them under the responsibilist norm for being ignorant of that data.

culpable for the fact that their beliefs fail to fit the evidence. Second, we could hold that only the publically available data of which the agent is in fact aware counts at all when it comes to determining whether or not the agent's belief fits the evidence.

The second option is preferable. Suppose the publically available data comprises evidence set E_1 , and that E_2 is a subset of E_1 consisting of the publically available data of which the agent is actually aware. Suppose further that E_1 fits some proposition P , while E_2 fits $\sim P$. What should the agent believe? She should believe $\sim P$. That is what fits the evidence – *her* evidence, which is the only evidence that could be relevant to evaluating her beliefs. If she believed P , she would be believing in a way that did not fit her evidence. It is true that we might wish that she were better informed, and in being better informed believed P , but when it comes to evaluating her performance in the situation in which she in fact finds herself, we can only say that she ought to believe $\sim P$. So it is simply wrong to say that an agent *should* believe in accordance with all the publically available data, even if she does not have relevant parts of that data, and even if we allow that she may not be culpable for failing to do so.

Publically Available Data Plus Awareness

So, suppose that when we try to determine whether someone else's belief fits the evidence, the evidence that we need to be looking at is the publically available data of which the agent is in fact aware. Here too, the task of determining what the other person's evidence is would be fairly straightforward. There are plenty of things we can unproblematically assume that other people know – things that are common knowledge and empirically obvious features of the other person's surroundings, for instance. And we can pretty confidently assume that people are aware of certain kinds of publically

available data based on what we know of their education and background. Yetter-Chappell, for instance, would be quite reasonable in assuming that his fellow academic philosophers are familiar with the arguments he takes to constitute the evidence relevant to theism. And of course, if we are in any doubt as to whether the agent is aware of a particular piece of the publically available data, we can in many cases ask her. Of course, that last option will not be available in cases in which the other person is, for whatever reason, beyond the reach of our communication. But if evidence is just the publically available data of which the agent is aware, we will very often be in a good position to know what the agent's evidence is.

But even if we know an agent's beliefs and know what publically available data she is aware of, that will not be enough to tell us whether her belief matches her evidence.

That is because we can imagine cases in which, in addition to all of the publically available data of which a person is aware, there also exists some spurious piece of publically available data the existence of which he falsely believes in. Suppose, for example, that he simply misheard the testimony of some reliable expert, such that he believes that she asserted that P when in fact she asserted that $\sim P$.

Since the expert did not assert that P , that the expert asserted that P is not part of the publically available data of which our agent is aware. He simply wrongly believes it to be. So suppose that the actual publically available data of which the agent is aware supports the conclusion that $\sim P$, but that said data along with the expert's assertion that P , had she in fact made such an assertion, would support the conclusion that P .

For the same reasons given above, it is wrong to say that he should believe that $\sim P$ under these circumstances. His own understanding of what the evidence is demands that he believe that P . How could he rationally do otherwise?

So it is not the case that we should always believe in accordance with the publically available data of which we are aware. So the evidence that we blame people for not believing in accordance with cannot be the publically available evidence of which they are aware.

These considerations might push us in the direction of saying that the relevant sense of evidence is whatever the subject *believes* to be the publically available data. If that were the case, we could still often have a fairly good idea of what the subject's evidence is. We can very often make reasonable assumptions about what someone believes the publically available data to be, and in many cases in which we have a pressing need to decide whether someone is blameworthy, we will be in a position to ask them what they believe the publically available data is.

So if the relevant sense of evidence is whatever the subject believes to be the publically available data, we will often be in a good position to make judgments about blameworthiness under the evidentialist norms.

But that cannot be the relevant sense of evidence either. To see why it cannot be, we must consider the importance of private evidence in our epistemic lives.

The Importance of Private Evidence

Private evidence – evidence that a subject cannot make known to others but which nevertheless plays a significant role in what she ought to believe under the evidentialist norm -- exists and matters. Return to our trusting husband. As the case has been

described, he has a set of publically available data that seems to indicate that his wife is having an affair. He also trusts his wife, and does not believe that she is having an affair. I argued above that we need not see anything epistemically untoward in this situation, because there is no reason not to treat his trust in his wife as part of the evidence he has to bring to bear on the question of her fidelity.

There is some sense, of course, in which his trust could be publically available data. After all, he is quite capable of saying “I trust my wife,” and we as third-party observers are quite capable of taking this fact into account as we decide what we think the evidence supports. He could reveal facts about her past behavior that seems to bear upon her trustworthiness. That is all publically available.

But while mere fact of the husband’s trust is publically available, its evidential force is not. We know that he trusts his wife, but why should that trust outweigh all the other evidence that seems to point to her infidelity? It is quite possible that there are good answers to that question that the man is entirely incapable of articulating. “I just can’t imagine her doing something like that,” he says. What he can or can’t imagine her doing is a result of a lifetime of experiences with her far too numerous to articulate or even, for that matter, recall. We cannot examine those experiences, so we cannot know whether his inability to imagine her having an affair is reasonable or a stubborn and blameworthy refusal to face the truth.

Or imagine a philosopher who is perfectly familiar with the theistic and atheistic arguments that Yetter-Chappell mentions, and agrees with his assessment of their relative merits. Nevertheless she remains a theist, on the basis of experiences she has had that were as of God talking to her.

Leave aside the question of whether there are defeaters for such experiences that should cause anyone aware of them to utterly discount the evidential value of said experiences. If there are any, assume that our philosopher reasonably, if falsely, believes those defeaters defeated.

It should be uncontroversial that experiences as of P constitute at least some evidence, however slight and however defeasible, for P. So our philosopher's religious experience should be included in the evidence she brings to bear on the question of God's existence.

But again, that evidence is not public. If we asked her why she believed in God despite the fact that she believes that the atheistic arguments are much stronger than the theistic arguments, she might say "I had an experience as of God speaking to me." But then the publically available evidence would be that she had reported such an experience, not the experience itself. She would be unable to communicate the qualitative character of the experience to us. In particular, she would not be able to show us the features of that experience that lead her to take it as sufficient reason to believe in God despite the acknowledged strength of the atheistic arguments. But without access to those features of the experience, we are not in a position to know whether an evidence set that includes that experience fits belief or disbelief in God.

Against private evidence? Yetter-Chappell objects to the claim that experiences of this sort can serve as private evidence.⁵ When S reports such an experience to me, he says, I come to believe (assuming I trust S) the proposition <S had such-and-such an experience.> That S had such-and-such an experience becomes evidence for me. But the

⁵ Richard Yetter-Chappell, "Experience and Testimony," <http://www.philosophyetc.net/2006/10/experience-and-testimony.html> October 16, 2006

evidence that S gains from the experience is exactly the same – that S had such-and-such an experience. So even with S's religious experience, S and I have the same evidence with respect to the existence of God. If that evidence doesn't fit belief in God for me, then neither does it do so for S.

The defender of private evidence, Yetter-Chappell says, has to claim that S's experience itself, rather than the proposition that S had the experience, serves as evidence for S. Only then could S and I have different evidence. But that, he says, "sounds bizarre."

It does not sound bizarre to me. There are, however, contemporary views of evidence that would rule out the possibility of S's experience itself being evidence. One such is view of Williamson (2002), which argues that a subject's evidence is the set of propositions that that subject knows. If that is the case, then of course nothing that is not a proposition is evidence, and therefore private, non-propositional experiences are not evidence.

We are not interested in everything that might be reasonably referred to by the word "evidence," and there is no compelling *prima facie* reason to suppose that everything that could reasonably be referred to by that term can be subject to a single unified analysis. All we have to ask, then, is whether the set of things we can be properly epistemically blamed for not believing in accordance with consists only of propositions.

Here's an argument that it does. Propositions are the only things that we can properly be epistemically blamed for failing to believe in accordance with because they are the only things we *can* believe in accordance with. Beliefs, after all, have propositional content, and that content is what determines whether or not they fit the

evidence. When we talk about beliefs fitting the evidence, we are talking about propositions fitting the evidence.

It is fairly easy to understand what it means for a proposition to fit the evidence if the evidence consists only of other propositions. There is the minimal and unimpressive kind of fitting in which the proposition believed is logically consistent with the propositions that make up the evidence. There is the maximal and impressive kind of fitting in which the proposition believed is logically entailed by the propositions that make up the evidence. And in between those extremes, the proposition believed can enjoy varying degrees of probability, given the propositions that make up the evidence, according to Bayes' theorem or some other well-defined system of inductive logic. So we have a fairly good grip on how beliefs could fit evidence if evidence consists entirely of propositions, and thus how someone might be blameworthy for not believing in a way that fit a purely propositional set of evidence.

It is, however, difficult to see how a believed proposition could be in any sort of epistemically relevant fitting relation with an experience. Propositions do not have the sort of inferential relationships with experiences that they have with one another. So, when it comes to evidence as that which we can be epistemically blamed for not believing in accordance with, propositional views seem to have an advantage insofar as it can offer us a good account of what that means.

But we need not settle that issue, since Yetter-Chappell is mistaken when he says that the defender of private evidence must say that experiences, as opposed to propositions about experiences, can be evidence.

He is mistaken because when S tells me that he has had a religious experience but is unable to communicate it in any kind of informative detail, the proposition that I come to believe on the basis of his testimony is not the same as the proposition that he comes to believe on the basis of his experience. “Such-and-such an experience” refer to different things in each of our beliefs. In the proposition he believes, it refers to the experience itself, in all its qualitative glory. In the proposition that I believe, it refers to something about which I know almost nothing other than the fact that S took it to have something to do with God.

You might object here that the experience about which I know almost nothing other than the fact that S took it to have something to do with God *just is* the experience that S had in all its qualitative glory, so the two propositions are in fact the same. But if that’s the case then we can believe propositions under different descriptions, and the differences in those descriptions can be evidentially relevant. If I believe that Superman is going on vacation in Sydney, that gives me evidence that the crime rate will increase in Metropolis and decrease in Sydney. Believing that Clark Kent is going on vacation in Sydney has no such evidential force in the absence of knowledge that he is Superman. So even if the proposition that I believe as a result of S’s testimony and the proposition that S believes as a result of his experience are the same proposition because they refer to the same object, the differences in the descriptions under which we believe those propositions can result in us having different evidence.

So, private evidence is real and important. A judgment about whether an agent’s belief fits the evidence that considers as evidence only the publically available data of which the agent is aware is inadequate.

Seeming Evidentialism

What, then, is evidence, for the purpose of blame under the evidentialist norm? What is it that we can be blamed for not believing in accordance with? I think that the best answer to this question is the view of evidence contained in what Earl Conee calls “seeming evidentialism.” On that view, “evidence is supplied by seeming truth. [...] [S]omeone’s evidence about a proposition includes all that seems to the person to bear on the truth of the proposition.⁶” If something seems to me to indicate that P, or if it seems to me on the basis of something that P, then that thing is evidence for P.

This is a broad view of evidence that can accommodate both propositional and experiential evidence. A proposition might seem to me to indicate the truth of some further proposition, insofar as I apprehend the inferential relationships between them. Some proposition might seem true to me on the basis of some experience. Since both the proposition and the experience play this role, they both count as evidence.

Why should we accept this view of evidence as the one that matters for the evidentialist norm? To answer that question, we answer two other questions. First, why should we accept the claim that nothing that does not seem to a subject S to bear upon the truth of proposition P is something that S can be epistemically blamed for not believing in accordance with when it comes to her beliefs concerning P? Second, why think that everything that seems to S to bear upon the truth of P is something that she can be properly blamed for not believing in accordance with when it comes to her beliefs concerning P?

⁶ Conee 2004, 15.

To take the first question, it would just be bizarre to blame anyone for not believing in accordance with something that they take to have no bearing on the truth or falsity of the proposition under consideration. How could they? Or at least, how could they see any reason to?

To take the second question, to adopt an attitude towards a claim is to judge the relative weight of all of one's various reasons to adopt various attitudes towards it. To adopt an attitude towards a claim without actually taking into consideration everything one takes to bear upon the truth of that claim would run counter to the purpose of forming an attitude.⁷

Can the evidentialist norm be violated? So, once we get clear on what counts as evidence for the purpose of the evidentialist norm, we see that it is very difficult to know what someone else's evidence is, and thus difficult to know whether they are culpably failed to believe in accordance with their evidence. This fact should introduce a bit of caution into our ascriptions of blame under the evidentialist norm.

But these conclusions give us another reason to be cautious about ascribing blame under the evidentialist norms as well. If the account of the relevant notion of evidence given here is correct, we must ask whether it is actually possible to violate the evidentialist norm. After all, what sort of creature does not believe on the basis of how things seem to him or her to be? Even if there is enough conceptual space between seemings and beliefs that it is possible for *S* to believe that *P* even though it does not on balance seem to *S* that *P*, or for it to seem to *S* that *P* without *S* believing that *P*, any such cases are going to be instances of massive malfunction rather than what we would take to

⁷ I am indebted to Trent Dougherty, in personal conversation, for this point.

be ordinary epistemic misbehavior.⁸ So we should expect violations of the evidentialist norm to be, at most, a rarity. This worry strengthens the note of caution sounded above about ascriptions of blame under the evidentialist norm.

The Responsibilist Norm

So, epistemic blame under the evidentialist norm is problematic. But what about blame under the responsibilist norm? That, at least, should not be a problem. After all, even if people's evidence is private, their behavior is public. We can have a pretty good idea of what kind of inquiry people have done. We can ask people to reveal the sources they have consulted, the lines of reasoning that they have consciously pursued, and so on. Surely that should not be a problem.

It is plausible enough that all of the things that a person could be required to do under the responsibilist norm are things that we will not have much trouble figuring out whether she has done or not. But that is not enough to make it the case that we can easily know whether someone has violated the responsibilist norm. We must also be able to easily know what someone's obligations under the responsibilist norms are to begin with.

But what someone's obligations under the responsibilist norms are depend, in part, on the evidence that they have. Take, for instance, the norm requiring diligent and thorough inquiry. Exactly how much effort I should think I need to expend to gather more evidence on a question depends on what I think of my current collection of evidence on that question. I think, for instance, that my collection of evidence on the question of whether I have hands is quite good. It supports the proposition that I have hands very strongly. In addition, it covers the *kinds* of available evidence rather thoroughly. If I

⁸ Again, I thank Trent Dougherty for help with this point in personal conversation.

continue inquiring into the question, the only pertinent additional evidence I'm likely to acquire will consist of more sensory experiences that are more or less the same as all my other relevant sensory experiences. So I have no reason at all to expect further evidence to come along that might cause me to no longer believe that I have hands. I have no reason to expect new evidence to come along that would cause me to revise my attitude downwards, and even if it did, my other evidence supports my having hands so strongly that it would take quite a lot of downward revision before I actually stopped believing that I have hands. Under those circumstances, nobody would say I was doing anything epistemically blameworthy by not enquiring any further into the question of whether I have hands.

Contrast this with the case of a religious person who believes in God on the basis of the testimony of her parents, neighbors, and pastor. While she believes these people to be fairly trustworthy on a wide range of topics, she is quite aware that it is possible that they are mistaken in this case. She is thus quite aware that her evidence that God exists is, at least, not overwhelming. Suppose she then becomes aware that other kinds of evidence on the question are available – that there exists a long tradition of philosophical reflection on the question of God's existence. It seems that she should think that there is an avenue of inquiry open to her that might cause her to cease to believe in God. After all, she thinks that her current evidence is not overwhelmingly strong, and she thinks that there's a class of evidence out there that for all she knows contains evidence that would cause her to revise her belief in God downward. It is at least plausible to say that she would be epistemically blameworthy for not pursuing this line of inquiry.

The fact that I do not have an obligation to enquire further and our hypothetical religious person does is a function of our evidence. My evidence is such that further enquiry is likely to be useless, while hers is such that it is likely to be useful. So our limited access to what evidence other people have also limits our access to what their obligations are under the responsibilist norm.

In light of these results, how should we proceed with the practice of epistemic blame? I will take up that question in the next chapter.

CHAPTER FIVE

Epistemic Blame and Charity

Introduction

So we have learned that, for at least a very large range of cases in which we might naturally be inclined to judge that a subject has culpably violated an epistemic norm, there remains the possibility that the subject has some evidence that we do not have access to and the presence of which means that the subject has in fact violated no epistemic norm. How seriously should we take such a possibility? Should we simply always, in the name of charity, assume that the subject is in possession of such evidence, or at least keep the possibility that she is sufficiently open to prevent us from ascribing epistemic blame? If so, what are we to do with the important functions epistemic blame serves, and with the seemingly obvious claim that people are sometimes epistemically blameworthy?

When we try to figure out what someone's evidence is for the purposes of determining whether they are epistemically blameworthy, we are attempting a kind of interpretation of that person. We interpret others in various ways all the time. We interpret people's beliefs, or assertions in languages partially or entirely unfamiliar to us, on the basis of their behavior, presumed desires, and so forth. We construct precise and formal accounts of people's reasoning on the basis of their unsophisticated reports of it.

In all such cases, we operate according to various principles of interpretive charity. Put very roughly, we interpret people's assertions, beliefs, desires, etc. such that their behavior and assertions make sense. If we watch Smith go to her refrigerator, get

out a beer, open it, and drink it, that generates for us the very strong presumption that Smith's desire was to have a beer. Alternate interpretations, such as ones on which Smith desires to avoid beer or is indifferent towards it, render her behavior inexplicable and are thus to be avoided.

Discussions of interpretive charity appear in countless texts on introductory logic and critical thinking, as well as more formally in Quine's work on radical translation and Davidson's work on radical interpretation. None of these discussions deal directly with our problem, which is the interpretation of a subject's evidential state. So we will have to see how the principles of charity contained in these discussions can be adapted to our problem. We will then see if these principles do justice to our belief that people are sometimes epistemically blameworthy. Finally, we will ask what reason we have to interpret the evidential states of others charitably.

Why be Charitable?

There are at least three kinds of reasons we might have for interpreting others charitably. I will call these alethic reasons, epistemic reasons, and moral reasons.

Alethic Reasons

We have an alethic reason to interpret other people charitably if, within the domain being interpreted, charitable interpretations are more likely to be true than uncharitable ones. Take the case in which Smith goes to the refrigerator, gets a beer, and drinks it. Interpretive charity tells us that we should interpret her desires such that they include the desire for a beer, because that is the interpretation on which Smith's desires are related to her actions in a rational way. The fact that on that interpretation Smith's

desires are related to her actions in a rational way is relevant because Smith's desires are extremely likely to be related to her actions in a rational way. So the interpretation of Smith's desires on which they are so related is, *prima facie*, much more likely to be true than any alternative. So we have an alethic reason to interpret Smith charitably here simply because that interpretation is likely to be true.

Do we have alethic reasons to interpret others' evidence charitably, i.e. such that they are epistemically blameless? In other words, are interpretations of other people's evidence on which they are epistemically blameless more likely to be true than interpretations on which they are not? That depends on what we think about the relative frequency of epistemic blameworthiness in the relevant cases.

One might reason as follows: The vast majority of the beliefs people hold are perfectly proper. Perceptual beliefs and memorial beliefs are normally unproblematic, and testimonial beliefs very often are as well. Evidence for most of our beliefs more or less falls into our laps. The evidentialist norm is easy to satisfy, and the bar to be cleared to satisfy the responsibilist norm simply is not very high. So we should expect the large majority of people's beliefs to be epistemically blameless. So interpretations of others' evidence on which they are epistemically blameless are more likely to be true than alternatives, and we therefore have alethic reasons to adopt them.

The problem with that line of reasoning is that we don't need a principle of charity to deal with the large range of easy cases in which someone is clearly epistemically blameless. We need them for the hard cases in which the question of epistemic blame arises in the first place. It is only when someone believes something that goes against our expectations of what a person inquiring responsibly and following her

evidence would believe that there is any need to make a judgment about epistemic blameworthiness.

So we only have alethic reasons to interpret others' evidence charitably if we think that blameless belief is more common than blameworthy belief in those hard cases. That is a more difficult case to make. I will return to it later in this chapter.

Epistemic Reasons

We have an epistemic reason to interpret someone charitably if charitable interpretations within that domain are, regardless of whether they themselves are more likely to be true, more likely than their alternatives to help us get to the truth. For an example of this sort of reason for charity, consider the charitable reconstructions of other people's arguments. Suppose that Smith has given an vague and unrefined presentation of her reasons for believing that P. The language she uses would equally support the formalizing of her reasoning into either one of two arguments, one of which is clearly invalid or depends on wildly implausible premises, and one of which is valid and relies on plausible premises.

Now, depending on what we think of Smith, we may have alethic reasons to interpret Smith charitably. We may say to ourselves "Smith is a smart and thoughtful person, so it's much more likely that her reasoning is best captured by the better of the two arguments." But perhaps we don't know whether Smith is a smart and thoughtful person, or even believe that she is not. In that case, we might not have a reason to think that the charitable interpretation of Smith's reasoning is more likely to be the true one.

Even so, we would still have a reason for preferring the charitable interpretation of Smith's reasoning. If we interpret Smith's reasoning uncharitably, all we have is an

uninterestingly bad argument that is of no use to anyone. But if we interpret Smith's reasoning by formalizing it into an argument that is valid and has plausible premises, then we might learn either that P or that at least one of the argument's plausible premises is false. That would be interesting and worthwhile. It would help us epistemically. Of course, we would not strictly need to ascribe the better argument to Smith in order to enjoy this benefit. We would simply need to become aware of it. But we are much more likely to become aware of the better argument if we are making a conscious effort to interpret Smith's reasoning as charitably as possible. So we have an epistemic reason to interpret the reasoning of others charitably in that doing so will help us in our overall goal of getting to the truth.

What epistemic reasons might we have for charitably interpreting other people's evidence? You might think that, just as interpreting reasoning charitably helps us get to the truth by helping us come upon good arguments, interpreting evidence charitably helps us get to the truth by helping us come upon good evidence. But that is not the case. The good arguments that we come upon by charitably interpreting people's reasoning are, simply by virtue of being good arguments, epistemically useful regardless of whether they accurately represent anyone's actual reasoning. The same cannot be said in the case of the interpretation of evidence.

For one thing, in interpreting another's evidence we are generally not reconstructing some specific and well-defined piece of evidence. We are not saying to ourselves, "Well, perhaps he had an experience of such-and-such a character, and that renders his present belief epistemically blameless." If we could do that, the experience would be communicable to us and the need for such guesswork would be eliminated. We

are, rather, hypothesizing the existence of *some* sort of evidence that renders the subject blameless. That's of no value when it comes to helping us get to the truth.

For another thing, merely hypothesized evidence has no evidential force. An argument someone *might* have intended to make is still an argument, and can do all the work that arguments can do for us regardless of whether anyone actually intended to make it. But possible experiences, along with other kinds of possible evidence, do nothing at all to indicate the truth of anything unless they are actually had by someone.

Moral Reasons

Finally, we can have moral reasons for interpreting others charitably. Doing so serves as a safeguard against rash judgment and pride. It helps us love our neighbor.

It seems clear that we have moral reasons to be charitable in our interpretation of other people's evidential states. If epistemic blame is a species of moral blame, then that reason is all the greater, since moral evaluation done the wrong way is at greater risk for the evils mentioned above.

But it seems that we could avoid these moral evils and achieve these moral goods just as easily by simply withholding any judgment at all as to whether someone is blameworthy.

Why not Withhold?

One might wonder whether we really do need any kind of principle of charity at all. After all, we are dealing with cases in which we do not have sufficient evidence to determine whether we should believe or disbelieve the proposition that some subject is epistemically blameworthy. The normally recommended response in such situations is

not to turn to principles to guide our guesswork. It is simply to refrain from either believing or disbelieving the proposition in question.

In many cases we have no pressing need to adopt any attitude at all towards the question of someone else's epistemic blameworthiness, and can and probably should refrain from doing so. But there are some cases in which we need such judgments to guide actions. Judgments of blame are often not mere trivia. We need them in order to navigate our social setting. Judges need to determine what a defendant ought to have believed in order to determine her culpability. I need to determine whether someone is acting in good epistemic faith when it comes to some point of contention between us so I can decide whether it is worthwhile to continue to engage him on the matter.

Also, if our principle of charity is motivated by alethic concerns, then we are not using principles of charity in order to make judgments about a subject's blameworthiness or lack thereof in the absence of actual evidence on the matter. An alethically-motivated principle of charity would constitute evidence in favor of a particular judgment about a subject's blameworthiness.

Principles of Charity

Here is a rough and ready principle of charity for epistemic blame: resolve doubts about a subject's blameworthiness in favor of their blamelessness. In other words, interpret people as blameless. In the remainder of this chapter, I will turn to work by Quine and Foley to defend and refine this principle.

Quine

In his discussion of radical translation in *Word and Object*, W.V.O Quine discusses the translation of a language's logical connectives.¹ For truth-functional connectives this is not a difficult task. A bit of language the addition of which turns every sentence that a native speaker of the language would assent to into one she would reject and every sentence she would reject into one she would assent to is a negation. A bit of language joining to statements to make a compound statement that a native speaker would assent to if and only if she assents to both of the atomic statements is a conjunction. And so on.

Suppose, though, that once we have decided upon translations for logical connectives, we find native speakers of the language we are trying to translate assenting to statements that are, according to our translations for the connectives, contradictory. So, for instance, speakers will assent to statements that use what we take to be the language's terms for negation and conjunction to make a statement of the form 'P and not P.' One option in such a case would be to go ahead and impute a contradiction to the speaker. Perhaps logical consistency simply is not something the speakers of the language prize, or even understand.

It would be better, Quine says, to alter the translation to preserve the logical coherence of the speaker's assertions. This fits with the way we interpret English. To use his example, if we ask an English speaker whether some statement is true or false and she responds "Yes and no," we do not interpret her as asserting the conjunction of that

¹ W. V. O Quine, *Word and Object* (Cambridge, MA: MIT Press, 1960) Pp 57-61

statement and its negation. We assume that she means that the statement can be understood in two ways, and is true on one understanding and false on the other.

Why do we make that assumption? We do it because to do otherwise would be to impute to the speaker an improbable degree of cognitive dysfunction. This suggests a general principle: “one’s interlocutor’s silliness, beyond a certain point, is less likely than bad translation – or, in the domestic case, linguistic divergence.”² Simply put, it is more likely that someone is using language in a way that we do not expect than that they are so daft as to affirm an explicit contradiction. This fact should govern the way we interpret other people’s use of language, whether it is our own language or one we are attempting to translate. We should recognize that “assertions startlingly false on the face of them are likely to turn on hidden differences of language.”³

Note that the reasons for charity here are clearly alethic. We should interpret language such that its users are abiding by basic laws of logic because it is extremely likely that they are.

How can this discussion be applied to our question? Adapting Quine’s maxims to the question of the interpretation of evidence would yield something like *beliefs startlingly blameworthy on the face of them are likely to turn on hidden evidence* and *one’s interlocutor’s blameworthiness, beyond a certain point, is less likely than hidden evidence*. These claims, if true, would provide alethic reasons for charitable interpretation of other people’s evidence.

² *Ibid.* 59

³ *Ibid.* 59

The unlikelihood of large norm violations. An advantage of these maxims is that they suggest to us something more detailed than “try to interpret other people’s evidence charitably.” Both maxims suggest that the likelihood of someone being rendered blameless by hidden evidence tracks with the degree of their belief’s apparent blameworthiness. The more wildly someone seems to have violated a norm, the more likely it is that they have not.

That suggestion, while it enjoys the advantage of offering more detailed guidance, might seem paradoxical. The more obvious it seems to us that someone is blameworthy, the less likely that they are. An odd suggestion indeed.

But the maxims do not commit us to anything quite so paradoxical as that. The claim is not that the more obvious it seems to us that someone has violated a norm, the less likely they have. The claim is that the greater the margin by which they seem to have failed to satisfy the norm, the less likely that they have. That is a much more plausible claim.

To see its plausibility, consider an example involving social norms. Suppose you having a conversation with Smith and Jones, two strangers you have just met at a party. In the course of casual conversation, Smith asks Jones what his salary is. This is a rather rude question. It violates social norms of politeness. And it seems, on the face of it, rather obvious that it does so. You would likely conclude from this exchange that Smith needs a few reminders about the niceties of social interaction.

Now suppose instead that, without provocation, Smith graphically insults Jones and every member of Jones’s family and ends his tirade by spitting in Jones’s face. This, to, is a violation of social norms of politeness. As in the first case, it seems like an

obvious violation of those norms. But this case is different in that it involves a much more significant violation of the norms of politeness. Smith has apparently failed to be polite by a much wider margin than in the first case.

In the first case, you would likely conclude that Smith is simply a bit rude, or at least had behaved rudely in this isolated case. But in the second case, you would quite possibly wonder if there is some alternate explanation for Smith's behavior. People are sometimes rude, after all, but it's rare that they are *that* rude. So you might wonder if Smith has some sort of condition that causes him to have uncontrollable outbursts for which he cannot be blamed. Or perhaps Smith and Jones are playing a joke on you. Or perhaps they are actors providing the evening's entertainment, and what you thought was your conversation with them was actually their attempt to act out their scene despite your frequent interruptions. You might eventually reject such alternate explanations and conclude that Smith is simply a man of unique and extreme rudeness. But it at least occurred to you to wonder about alternate explanations in the second case when it did not occur to you to wonder about them in the first case, precisely because in the second case Smith seems to be falling short of the norms of politeness by an improbably large margin.

So in at least some realms, it is perfectly sensible to take the large margin by which someone seems to you to be failing to fulfill a norm as evidence that they are in fact not violating a norm at all.

But it's sensible in the case described above because we have good empirical reasons for thinking that people don't often miss the mark by a wide margin when it comes to social norms, or at least that they're less likely to miss it by a large margin than by a small. We observe people behaving in social situations all the time. We know that

massive violations of the social norms are rare and noteworthy, not a standard feature of social life. It's that knowledge that causes us to suspect that what appears to be a massive violation is in fact something else.

The question, then, is whether we have reason to think that the same is true either particularly of epistemic norms or of norms in general (and thus of epistemic norms.) In the next section I will argue that this is the case.

Charity, Self-Trust, and Trust in Others

Making the case that we have good reason to think that serious violations of the epistemic norms are relatively rare requires dealing with some of the recent literature on self-trust. In what follows I will argue that the epistemic trust that we must place in ourselves commits us to the position that people generally abide by epistemic norms.

Self-Trust and Skepticism

There is now widespread agreement that the classical foundationalist project of providing a complete system of knowledge based on indubitable, skeptic-satisfying foundations has no prospects for success. This is because the attempt to provide such foundations is doomed to be circular. It is doomed to be circular because it must appeal, at some point, to certain first principles of reasoning and to basic human cognitive faculties. But the truth of such principles and the reliability of such faculties are part of what a complete answer to the skeptic would require us to argue for. So any attempt to refute the skeptic will proceed from premises that, by its own lights, it must establish by argument. It would thus be uselessly circular.

The acceptance of this criticism of the foundationalist project does not, however, necessitate surrender to the skeptic. We do not have to choose between refuting the skeptic using premises she would accept and admitting that all of our normal cognitive lives are without foundation and thus irrational. We can instead attempt to justify our reliance on certain first principles and on our basic cognitive faculties in some way that does not involve conclusively establishing their reliability.

One way to do this is by appealing to the unavoidability of our basic faculties and first principles. A classic expression of this point is found in Thomas Reid:

"All reasoning must be from first principles; and for first principles no other reason can be given but this, that, by the constitution of our nature, we are under a necessity of assenting to them. Such principles are parts of our constitution, no less than the power of thinking: reason can neither make nor destroy them; nor can it do anything without them: it is like a telescope, which may help a man see farther who has eyes; but, without eyes, a telescope shews nothing at all. [...] How or when I got such first principles, upon which I build all my reasoning, I know not; for I had the before I can remember: but I am sure they are parts of my constitution, and that I cannot throw them off. [...] That our sensations of touch indicate something external, extended, figured, hard or soft, is not a deduction of reason, but a natural principle. The belief of it, and the very conception of it, are equally parts of our constitution. If we are deceived in it, we are deceived by Him that made us, and there is no remedy."⁴

Here the reason for believing that, for instance, that the deliverances of our sense of touch are indicative of an external world is not any argument that this is so, but the sheer impossibility of doing anything else. Believing in an external world on the basis of our sense of touch (among other things) is so fundamental to our cognitive nature that if doing so is an error, then we are unavoidably given to error. We cannot avoid that error simply by becoming skeptics, because we cannot become skeptics.⁵ If we are in fact

⁴ Thomas Reid. *Inquiry into the Human Mind*, ed. Derek R Brookes, (University Park, PA: Pennsylvania State University Press, 1997), 71-72

mistaken then our epistemic situation is tragic, but that does not make us any less reasonable in believing as we do, given that we could not do otherwise.

Richard Foley (2001) advances a related line of thought. For him, since there are no refutations of skepticism that do not beg the question available, “trust in one’s own intellectual faculties, procedures, and opinions is a part of any nonskeptical intellectual life.”⁶ I must place my trust in my own ability to get to the truth, and I am thus rational so to trust.

Self-Trust and Trust in Others

But if I am to place such trust in myself, Foley argues, it is incoherent for me to not likewise place such trust in others. There are multiple ways to arrive at this conclusion. The first simply appeals to the fact that I have no reason to believe that I am in a uniquely privileged epistemic situation. My cognitive equipment is roughly the same, in both kind and quality, as most other people’s. I do not have any senses that others lack to allow me to gather information that they have no access to. I am not vastly more intelligent than them such that I should think that I am far better at making inferences than they are. And I have been shaped by roughly the same cultural and environmental forces that shaped them. In short, it seems implausible to say that there is anything terribly special about me, epistemically speaking. But extending trust to my own faculties without extending trust to the faculties of others would be to treat myself as epistemically special. So I ought not do that.

⁵ As Reid points out, even the alleged skeptics to whom he is responding make free use of certain basic principles, such as the existence of sensations and ideas.

⁶ Richard Foley, *Intellectual Trust in Oneself and Others*, (Cambridge: Cambridge University Press, 2001), 99

That is a powerful reason to extend the same trust to others that I extend to myself. But as an argument for extending trust to others, it does have some weaknesses. For starters, it is conceivable that I might think that I am in fact epistemically special. It is even conceivable that, in some circumstances, I may do so blamelessly. Or I might reason that my goal is to minimize error, which means extending trust to as few people as possible. In such a case I might trust myself since otherwise I could not function at all, while declining to extend trust to others on the ground that I can get by without doing so and not doing so minimizes my chances of error.

What is needed, then, is an argument that to extend trust to myself but not others is downright incoherent -- that not trusting others undermines the trust I place in myself. Foley finds such an argument in the social embeddedness of our cognitive practices.⁷ The self in which I am forced to place my trust does not exist in cognitive isolation. I have not created the way that I look at the world from scratch. From childhood on I have been inundated with the beliefs of others. My parents told me all manner of things. My broader cultural context supplies me with countless presuppositions and concepts for ordering my understanding of the world. Influences like this form the background against which I form beliefs about the world. I cannot simply separate these influences from some unadulterated cognitive self in whom I must place my trust. They are simply part of what I bring to the table as an epistemic agent.

Not extending trust to others means not extending trust to those influences. Not extending trust to those influences means distrusting myself as an epistemic agent. So the self-trust that is necessary to avoid skepticism requires that I trust others as well.

⁷ *ibid* 102

Trust in Others and Epistemic Blame

How does this trust that we must extend to one another relate to our topic, which is whether we should believe that people are, in most cases, epistemically blameless? I see two ways of leveraging Foley's argument for trusting others into an argument for holding that others are, in general, epistemically blameless.

One involves a straightforward inference from reliability to blamelessness. We might argue that people cannot be generally reliable without being generally blameless. After all, to have incurred epistemic blame is to have ignored one or more of the norms meant to help us be reliable at getting to the truth. The blameworthy have failed to be responsible inquirers, or have failed to responsibly handle the evidence with which their inquiry has provided them. People such as that are not reliable at getting to the truth, at least not with regards to the matters about which they are blameworthy. So, since blamelessness is necessary for reliability, it must be at least as common as reliability. So any argument that I must trust that people are in general reliable serves equally well as an argument that I must trust that people are in general epistemically blameless.

The inference from reliability to blamelessness is not an ironclad one. We could imagine a world governed by an evil demon who, wanting to separate virtue from reward, arranges things such that epistemically blameworthy behaviors are very likely to result in true beliefs, and blameless behaviors very likely to result in false ones. Reliability in such a world would certainly not require blamelessness. And even in the actual world, we can imagine someone coming to some belief in a blameworthy way that nevertheless happens to be true. This person could be quite competent at forming good inferences based on the true but blameworthy belief. They would thus be quite reliable at getting to the truth with

respect to those inferences, but we would not call them blameless. So, while it seems right to take reliability as some sort of indicator of blamelessness, it is not clear that the connection between the two is strong enough for a general presumption of reliability to necessitate a general presumption of blamelessness.

Another strategy for extending Foley's results to our discussion involves making use of what he says about the connection between trusting others and trusting ourselves. If Foley is right, then I must make use of the input of others when it comes to my enquiry about the world. Doing that means trusting that they are generally reliable. Ought I do that without also trusting that they are generally epistemically blameless? Perhaps even more pressingly, *can* I do that without also trusting that they are generally epistemically blameless?

I do not believe that I can. Even if reliability does not entail blamelessness, it seems quite out of place to say that I trust someone to be generally reliable while denying that they are generally blameless, or even considering it to be an open question whether they are generally blameless. Reliability may not require blamelessness, but *trust* in reliability does require *trust* in blamelessness. To believe that someone is blameworthy or to be undecided on the matter is to believe in or be undecided on a rather large potential source of unreliability. In such a situation, saying that I trusted in the other's reliability would be very odd indeed. If you asked me whether I trusted Smith to drive me to the airport safely, and I said yes, but added that I did not trust him not to be drunk while doing so, you would rightly suspect that I was using the word "trust" in a perverse way. Likewise, if I say that I trust others to be reliable but do not trust them to be epistemically blameless, I am simply misusing the word "trust."

So, the trust I must have in myself means extending trust to others in a way that requires me to assume that they are usually epistemically blameless. This conclusion supports the claims about the frequency of epistemic blameworthiness that underlie the principles of charity advocated above.

BIBLIOGRAPHY

- Arneson, Richard. "The Smart Theory of Moral Responsibility and Desert." In *Desert and Justice*, edited by Sarena Olsaretti, 233-258. Oxford: Clarendon 2003.
- Clifford, W. K. *The Ethics of Belief and Other Essays*. Amherst, NY: Prometheus, 1999.
- Coates, D. Justin, and Neal A. Tognazzini. "The Nature and Ethics of Blame." *Philosophy Compass* 7, no 3 (2012): 197-207.
- Conee, Earl. "First Things First." In *Evidentialism: Essays in Epistemology*, edited by Earl Conee and Richard Feldman. 11-36 Oxford: Clarendon (2004)
- David, Marian. "Truth as the Primary Epistemic Goal: A Working Hypothesis." In *Contemporary Debates in Epistemology*, edited by Matthias Steup and Ernest Sosa, 296-312. Blackwell 2005
- De Neys, Wim. "Dual Processing in Reasoning Two systems but One Reasoner." *Psychological Science* vol 17 no 5 (2006) 428-433
- Dougherty, Trent. "Reducing Responsibility: An Evidentialist Account of Epistemic Blame." *European Journal of Philosophy* vol 20 no 4 (2012): 534-547
- Firth, Roderick. "Epistemic Merit, Intrinsic and Instrumental." *The American Philosophical Association Centennial Series* (2013): 5-18.
- Foley, Richard. *Intellectual Trust in Oneself and Others*. Cambridge: Cambridge University Press, 2001.
- Fumerton, Richard. "Epistemic Justification and Normativity." In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, edited by Matthias Steup, 49-58 Oxford: OUP, 2001.
- Haack, Susan. "The Ethics of Belief Reconsidered." In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*, edited by Matthias Steup, 21-30. Oxford: OUP, 2001.
- Hughes, Paul M., "Forgiveness", *The Stanford Encyclopedia of Philosophy* (Spring 2015 Edition), Edward N. Zalta (ed.),
URL <<http://plato.stanford.edu/archives/spr2015/entries/forgiveness/>>.

- Kvanvig, Jonathan L. "Truth and the Epistemic Goal." In *Contemporary Debates in Epistemology*, edited by Matthias Steup and Ernest Sosa, 285-296. Blackwell 2005.
- Locke, John. *An Essay Concerning Human Understanding*. Edited by Peter H. Nidditch. Oxford: Clarendon Press, 1975
- Malle, Bertram F., Steve Guglielmo, and Andrew E. Monroe. "A Theory of Blame." *Psychological Inquiry* vol 25 no 2 (2014) 147-186
- Millikan, Ruth Garrett. "In Defense of Proper Functions." *Philosophy of Science* vol. 56 no. 2. (1989): 288-302.
- Nottelmann, Nikolaj. *Blameworthy Belief: A Study in Epistemic Deontologism*. Vol. 338. Springer Science & Business Media, 2007.
- Nowell-Smith, P. H. *Ethics*. London: Penguin, 1954.
- Quine, W. V. O. *Word and Object*. Cambridge, MA: MIT Press 1960
- Reid, Thomas. *Inquiry into the Human Mind*. Edited by Derek R Brookes. University Park, PA: Pennsylvania State University Press, 1997.
- Scanlon, T. M. *Moral Dimensions: Permissibility, Meaning, and Blame*. Cambridge, MA: Harvard University Press, 2008.
- Sher, George. *In Praise of Blame*. Oxford: OUP, 2005.
- Smart, J. J. C "An Outline of a System of Utilitarian Ethics." In *Utilitarianism: For and Against*, Edited by J. J. C. Smart and Bernard Williams. Cambridge: Cambridge University Press. 1973
- Strawson, Peter Frederick. *Freedom and Resentment and Other Essays*. London: Methuen & Co, 1974.
- Tognazzini, Neal A. "Blameworthiness and the Affective Account of Blame." *Philosophia* vol. 41 no. 4 (2013)
- Williamson, Timothy. *Knowledge and its Limits*. Oxford: OUP, 2002
- Yetter-Chappell, Richard. "Experience and Testimony."
<http://www.philosophyetc.net/2006/10/experience-and-testimony.html>
 October 16, 2006

Yetter-Chappell, Richard. "Is Religious Belief Reasonable?"
<http://www.philosophyetc.net/2013/10/is-religious-belief-reasonable.html>
October 19, 2013

Zagzebski, Linda Trinkaus. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. Cambridge: Cambridge University Press, 1996.